



MIPS® Architecture for Programmers Volume II-B: microMIPS64™ Instruction Set

**Document Number: MD00594
Revision 6.05
December 15, 2016**

Copyright © 2016 Imagination Technologies LTD. and/or its Affiliated Group Companies. All rights reserved.

Public. This publication contains proprietary information which is subject to change without notice and is supplied ‘as is’, without any warranty of any kind.

Template: nB1.03, Built with tags: 2B ARCH MIPS64

MIPS® Architecture for Programmers Volume II-B: microMIPS64™ Instruction Set, Revision 6.05

Copyright © 2016 Imagination Technologies LTD. and/or its Affiliated Group Companies. All rights reserved.

Contents

Chapter 1: About This Book	15
1.1: Typographical Conventions	16
1.1.1: Italic Text	16
1.1.2: Bold Text	16
1.1.3: Courier Text	16
1.2: UNPREDICTABLE and UNDEFINED	16
1.2.1: UNPREDICTABLE	16
1.2.2: UNDEFINED	17
1.2.3: UNSTABLE	17
1.3: Special Symbols in Pseudocode Notation	17
1.4: Notation for Register Field Accessibility	20
1.5: For More Information	22
Chapter 2: Introduction	23
2.1: Default ISA Mode	23
2.2: Software Detection	24
2.3: Compliance and Subsetting	24
2.4: ISA Mode Switch	24
2.5: Branch and Jump Offsets	25
2.6: Coprocessor Unusable Behavior	25
2.7: Release 6 of the MIPS Architecture	26
Chapter 3: Guide to the Instruction Set	33
3.1: Understanding the Instruction Fields	33
3.1.1: Instruction Fields	35
3.1.2: Instruction Descriptive Name and Mnemonic	35
3.1.3: Format Field	35
3.1.4: Purpose Field	36
3.1.5: Description Field	36
3.1.6: Restrictions Field	37
3.1.7: Availability and Compatibility Fields	37
3.1.8: Operation Field	38
3.1.9: Exceptions Field	38
3.1.10: Programming Notes and Implementation Notes Fields	38
3.2: Operation Section Notation and Functions	39
3.2.1: Instruction Execution Ordering	39
3.2.2: Pseudocode Functions	39
3.3: Op and Function Subfield Notation	51
3.4: FPU Instructions	51
Chapter 4: Instruction Formats	53
4.1: Instruction Stream Organization and Endianness	56
Chapter 5: microMIPS Instruction Set	59
5.1: 16-Bit Category	59
5.1.1: Frequent MIPS64 Instructions	59

5.1.2: Frequent MIPS64 Instruction Sequences	62
5.1.3: Instruction-Specific Register Specifiers and Immediate Field Encodings	63
5.2: 16-bit Instruction Register Set	64
5.3: 32-Bit Category	66
5.3.1: New 32-bit instructions	66
5.4: Instructions Specific to microMIPS	68
ADDIUR1SP	69
ADDIUR2	70
ADDIUS5	72
ADDIUSP	74
ADDU16	76
AND16	77
ANDI16	78
BC16	79
BEQZC16	80
BNEZC16	81
BREAK16	82
JALRC16	83
JRCADDIUSP	85
JRC16	87
LBU16	88
LDM	90
LDP	92
LHU16	94
LI16	95
LWP	96
LW16	97
LWM32	98
LWM16	100
LWGP	102
LWSP	103
MOVE16	104
MOVEP	105
NOT16	107
OR16	108
SB16	109
SDBBP16	110
SDM	111
SDP	113
SH16	114
SLL16	115
SRL16	116
SUBU16	118
SW16	119
SWSP	120
SWM16	121
SWM32	123
SWP	125
XOR16	126
5.5: Recoded MIPS Instructions	127
ABS.fmt	128
ADD	129
ADD.fmt	130

ADDIU	131
ADDIUPC	132
ADDU	133
ALIGN DALIGN	134
ALUIPC	136
AND	137
ANDI	138
AUI DAUI DAHI DATI	139
AUIPC	141
BALC	142
BC1EQZC BC1NEZC	143
BC2EQZC BC2NEZC	145
B{LE,GE,GT,LT,EQ,NE}ZALC	147
B<cond>C	149
BC	152
BREAK	153
BITSWAP DBITSWAP	154
BOVC BNVC	156
CACHE	158
CACHEE	164
CEIL.L.fmt	170
CEIL.W.fmt	171
CFC1	172
CFC2	174
CLASS.fmt	175
CLO	177
CLZ	178
CMP.condn.fmt	179
COP2	184
CRC32B, CRC32H, CRC32W, CRC32D	186
CRC32CB, CRC32CH, CRC32CW, CRC32CD	189
CTC1	192
CTC2	195
CVT.D.fmt	196
CVT.L.fmt	197
CVT.S.fmt	198
CVT.W.fmt	199
DADD	200
DADDIU	201
DADDU	202
DCLO	203
DCLZ	204
DERET	205
DEXT	206
DEXTM	208
DEXTU	210
DI	212
DINS	213
DINSM	215
DINSU	217
DIV.fmt	219
DIV MOD DIVU MODU DDIV DMOD DDIVU DMODU	220
DMFC0	224

DMFC1	225
DMFC2	226
DMTC0	227
DMTC1	228
DMTC2	229
DROTR	230
DROTR32	231
DROTRV	232
DSBH	233
DSHD	234
DSLL	235
DSLL32	236
DSLLV	237
DSRA	238
DSRA32	239
DSRAV	240
DSRL	241
DSRL32	242
DSRLV	243
DSUB	244
DSUBU	245
DVP	246
EHB	249
EI	250
ERET	251
ERETNC	252
EVP	254
EXT	256
FLOOR.L.fmt	258
FLOOR.W.fmt	259
GINVI	260
GINVT	262
INS	265
JALRC	267
JALRC.HB	269
JIALC	272
JIC	274
LB	276
LBE	277
LBU	278
LBUE	279
LD	280
LDC1	281
LDC2	282
LDPC	283
LH	284
LHE	285
LHU	286
LHUE	287
LL	288
LLD	290
LLE	291
LLDP	294

LLWP.....	296
LLWPE	298
LSA DLSA	300
LUI	301
LW	302
LWC1	303
LWC2	304
LWE.....	305
LWPC.....	306
LWUPC	307
LWU	308
MADDF.fmt MSUBF.fmt	309
MAX.fmt MIN.fmt MAXA.fmt MINA.fmt.....	311
MFC0.....	315
MFC1.....	316
MFC2.....	317
MFHC0	318
MFHC1	320
MFHC2	321
MOV.fmt	322
MTC0.....	323
MTC1.....	324
MTC2.....	325
MTHC0	326
MTHC1	328
MTHC2	329
MUL MUH MULU MUHU DMUL DMUH DMULU DMUHU	330
MUL.fmt.....	333
NEG.fmt.....	334
NOP	335
NOR	336
OR	337
ORI	338
PAUSE	339
PREF	341
PREFE	345
RDHWR.....	349
RDPGPR	352
RECIP.fmt	353
RINT.fmt	354
ROTR	356
ROTRV.....	357
ROUND.L.fmt	358
ROUND.W.fmt.....	359
RSQRT.fmt.....	360
SB.....	361
SBE	362
SC	363
SCE	367
SCD.....	370
SCDP	373
SCWP.....	376
SCWPE	380

SD	383
SDBBP	384
SDC1	385
SDC2	386
SEB	387
SEH	388
SEL.fmt.....	390
SELEQZ SELNEZ	391
SELEQZ.fmt SELNEQZ.fmt	393
SH	395
SHE	396
SIGRIE	397
SLL	398
SLLV	399
SLT	400
SLTI	401
SLTIU	402
SLTU	403
SQRT.fmt	404
SRA	405
SRAV	406
SRL	407
SRLV	408
SSNOP	409
SUB	410
SUB.fmt	411
SUBU	412
SW	413
SWE	414
SWC1	415
SWC2	416
SYNC	417
SYNCI	422
SYSCALL	425
TEQ	426
TGE	427
TGEU	428
TLBINV	430
TLBINVF	433
TLBP	435
TLBR	436
TLBWI	438
TLBWR	440
TLT	442
TLTU	443
TNE	444
TRUNC.L.fmt.....	445
TRUNC.W.fmt	446
WAIT	447
WRPGPR	449
WSBH	450
XOR	451
XORI	452

Chapter 11: Opcode Map	453
11.1: Major Opcodes	453
11.2: Minor Opcodes	455
11.3: Floating Point Unit Instruction Format Encodings	465
Chapter 12: Compatibility	469
12.1: Assembly-Level Compatibility.....	469
12.2: ABI Compatibility	470
12.3: Branch and Jump Offsets	471
12.4: Relocation Types.....	471
12.5: Boot-up Code shared between microMIPS64 and MIPS64	471
12.6: Coprocessor Unusable Behavior.....	472
12.7: Other Issues Affecting Software and Compatibility	472
Appendix C: References	473
Appendix D: Revision History	475

List of Figures

Figure 3.1: Example of Instruction Description	34
Figure 3.2: Example of Instruction Fields	35
Figure 3.3: Example of Instruction Descriptive Name and Mnemonic	35
Figure 3.4: Example of Instruction Format	35
Figure 3.5: Example of Instruction Purpose	36
Figure 3.6: Example of Instruction Description	36
Figure 3.7: Example of Instruction Restrictions	37
Figure 3.8: Example of Instruction Operation	38
Figure 3.9: Example of Instruction Exception	38
Figure 3.10: Example of Instruction Programming Notes	39
Figure 3.11: COP_LW Pseudocode Function	39
Figure 3.12: COP_LD Pseudocode Function	40
Figure 3.13: COP_SW Pseudocode Function	40
Figure 3.14: COP_SD Pseudocode Function	40
Figure 3.15: CoprocessorOperation Pseudocode Function	41
Figure 3.16: MisalignedSupport Pseudocode Function	41
Figure 3.17: AddressTranslation Pseudocode Function	42
Figure 3.18: LoadMemory Pseudocode Function	42
Figure 3.19: StoreMemory Pseudocode Function	43
Figure 3.20: Prefetch Pseudocode Function	43
Figure 3.21: SyncOperation Pseudocode Function	44
Figure 3.22: ValueFPR Pseudocode Function	44
Figure 3.23: StoreFPR Pseudocode Function	45
Figure 3.24: CheckFPEException Pseudocode Function	46
Figure 3.25: FPConditionCode Pseudocode Function	46
Figure 3.26: SetFPConditionCode Pseudocode Function	47
Figure 3.27: sign_extend Pseudocode Functions	47
Figure 3.28: memory_address Pseudocode Function	48
Figure 3.29: Instruction Fetch Implicit memory_address Wrapping	49
Figure 3.30: AddressTranslation implicit memory_address Wrapping	49
Figure 3.31: SignalException Pseudocode Function	49
Figure 3.32: SignalDebugBreakpointException Pseudocode Function	49
Figure 3.33: SignalDebugModeBreakpointException Pseudocode Function	50
Figure 3.34: NullifyCurrentInstruction PseudoCode Function	50
Figure 3.35: NotWordValue Pseudocode Function	50
Figure 3.36: PolyMult Pseudocode Function	50
Figure 4.1: 16-Bit Instruction Formats	54
Figure 4.2: 32-Bit Instruction Formats	55
Figure 4.3: Immediate Fields within 32-Bit Instructions	55
Figure 5.1: ALIGN operation (32-bit)	135
Figure 5.2: DALIGN operation (64-bit)	135
Figure 5.3: Usage of Address Fields to Select Index and Way	158
Figure 5.4: Usage of Address Fields to Select Index and Way	164
Figure 5.5: Operation of the DEXT Instruction	206
Figure 5.6: Operation of the DEXTM Instruction	208
Figure 5.7: Operation of the DEXTU Instruction	210
Figure 5.8: Operation of the DINS Instruction	213

Figure 5.9: Operation of the DINSM Instruction	215
Figure 5.10: Operation of the DINSU Instruction	217
Figure 5.11: Operation of the EXT Instruction	256
Figure 5.12: Operation of the INS Instruction	265
Figure 11.1: Sample Bit Encoding Table	453

List of Tables

Table 1.1: Symbols Used in Instruction Operation Statements.....	17
Table 1.2: Read/Write Register Field Notation	20
Table 2.1: Instructions Added in Release 6	26
Table 2.2: Instructions Recoded in Release 6	28
Table 2.3: Instructions Removed in Release 6	29
Table 3.1: AccessLength Specifications for Loads/Stores	43
Table 4.1: microMIPS Opcode Formats.....	56
Table 5.1: 16-Bit Re-encoding of Frequent MIPS64 Instructions.....	60
Table 5.2: 16-Bit Re-encoding of Frequent MIPS64 Instruction Sequences.....	62
Table 5.3: Instruction-Specific Register Specifiers and Immediate Field Values	63
Table 5.4: 16-Bit Instruction General-Purpose Registers - \$2-\$7, \$16, \$17	64
Table 5.5: SB16, SH16, SW16 Source Registers - \$0, \$2-\$7, \$17	65
Table 5.6: 16-Bit Instruction Implicit General-Purpose Registers	66
Table 5.7: 16-Bit Instruction Special-Purpose Registers.....	66
Table 5.8: 32-bit Instructions introduced within microMIPS	66
Table 5.9: Encoded and Decoded Values of the Immediate Field.....	70
Table 5-1: Encoded and Decoded Values of Signed Immediate Field	72
Table 5.10: Encoded and Decoded Values of Immediate Field.....	74
Table 5-2: Encoded and Decoded Values of Immediate Field.....	78
Table 5.11: Offset Field Encoding Range -1, 0..14	88
Table 5.12: LI16 -1, 0..126 Immediate Field Encoding Range.....	95
Table 5.13: Encoded and Decoded Values of the Enc_Dest Field	105
Table 5.14: Encoded and Decoded Values of the Enc_rs and Enc_rt Fields	105
Table 5.15: Shift Amount Field Encoding.....	115
Table 5.16: Shift Amount Field Encoding.....	116
Table 5.17: Usage of Effective Address.....	158
Table 5.18: Encoding of Bits[17:16] of CACHE Instruction	159
Table 5.19: Encoding of Bits [20:18] of the CACHE Instruction	160
Table 5.20: Usage of Effective Address.....	164
Table 5.21: Encoding of Bits[22:21] of CACHEE Instruction	165
Table 5.22: Encoding of Bits [20:18] of the CACHEE Instruction	166
Table 5.1: Types of Global TLB Invalidates	262
Table 7.24: Special Cases for FP MAX, MIN, MAXA, MINA.....	313
Table 8.25: Values of <i>hint</i> Field for PREF Instruction	341
Table 8.26: Values of <i>hint</i> Field for PREFE Instruction.....	346
Table 8.27: RDHWR Register Register Numbers	349
Table 10.28: Encodings of the Bits[10:6] of the SYNC instruction; the STYPE Field.....	419
Table 11.1: Symbols Used in the Instruction Encoding Tables.....	454
Table 11.2: microMIPS64 Encoding of Major Opcode Field	455
Table 11.3: Legend for Minor Opcode Tables	456
Table 11.4: POOL32A Encoding of Minor Opcode Field	456
Table 11.5: POOL32P Encoding of Minor Extension Field	457
Table 11.6: POOL32Axf Encoding of Minor Opcode Extension Field.....	457
Table 11.7: POOL32F Encoding of Minor Opcode Field	459
Table 11.8: POOL32Fxf Encoding of Minor Opcode Extension Field.....	459
Table 11.9: POOL32B Encoding of Minor Opcode Field	460
Table 11.10: POOL32C Encoding of Minor Opcode Field	460

Table 11.11: LD-EVA Encoding of Minor Opcode Field.....	461
Table 11.12: ST-EVA Encoding of Minor Opcode Field.....	461
Table 11.13: POOL32I Encoding of Minor Opcode Field.....	461
Table 11.14: POOL32S Encoding of Minor Opcode Field	462
Table 11.15: POOL32Sxf Encoding of Minor Opcode Extension Field.....	463
Table 11.16: POOL16A Encoding of Minor Opcode Field	463
Table 11.17: POOL16B Encoding of Minor Opcode Field	464
Table 11.18: POOL16C Encoding of Minor Opcode Field	464
Table 11.19: POOL16D Encoding of Minor Opcode Field	464
Table 11.20: POOL16E Encoding of Minor Opcode Field	465
Table 11.21: PCREL Encoding of Minor Opcode Field.....	465
Table 11.22: Floating Point Unit Format Encodings - S, D, PS.....	465
Table 11.23: Floating Point Unit Format Encodings - S, D 1-bit	465
Table 11.24: Floating Point Unit Instruction Format Encodings - S, D 2-bits	466
Table 11.25: Floating Point Unit Format Encodings - S, W, L.....	466
Table 11.26: Floating Point Unit Format Encodings - D, W, L	466

About This Book

The MIPS® Architecture for Programmers Volume II-B: microMIPS64™ Instruction Set comes as part of a multi-volume set.

- Volume I-A describes conventions used throughout the document set, and provides an introduction to the MIPS64® Architecture
- Volume I-B describes conventions used throughout the document set, and provides an introduction to the microMIPS™ Architecture
- Volume II-A provides detailed descriptions of each instruction in the MIPS64® instruction set
- Volume II-B provides detailed descriptions of each instruction in the microMIPS64™ instruction set
- Volume III describes the MIPS64® and microMIPS64™ Privileged Resource Architecture which defines and governs the behavior of the privileged resources included in a MIPS® processor implementation
- Volume IV-a describes the MIPS16e™ Application-Specific Extension to the MIPS64® Architecture. Beginning with Release 3 of the Architecture, microMIPS is the preferred solution for smaller code size. Release 6 removes MIPS16e: MIPS16e cannot be implemented with Release 6.
- Volume IV-b describes the MDMX™ Application-Specific Extension to the MIPS64® Architecture and microMIPS64™. With Release 5 of the Architecture, MDMX is deprecated. MDMX and MSA can not be implemented at the same time. Release 6 removes MDMX: MDMX cannot be implemented with Release 6.
- Volume IV-c describes the MIPS-3D® Application-Specific Extension to the MIPS® Architecture. Release 6 removes MIPS-3D: MIPS-3D cannot be implemented with Release 6.
- Volume IV-d describes the SmartMIPS® Application-Specific Extension to the MIPS32® Architecture and the microMIPS32™ Architecture and is not applicable to the MIPS64® document set nor the microMIPS64™ document set. Release 6 removes SmartMIPS: SmartMIPS cannot be implemented with Release 6, neither MIPS32 Release 6 nor MIPS64 Release 6.
- Volume IV-e describes the MIPS® DSP Module to the MIPS® Architecture.
- Volume IV-f describes the MIPS® MT Module to the MIPS® Architecture
- Volume IV-h describes the MIPS® MCU Application-Specific Extension to the MIPS® Architecture
- Volume IV-i describes the MIPS® Virtualization Module to the MIPS® Architecture
- Volume IV-j describes the MIPS® SIMD Architecture Module to the MIPS® Architecture

1.1 Typographical Conventions

This section describes the use of *italic*, **bold** and `courier` fonts in this book.

1.1.1 Italic Text

- is used for *emphasis*
- is used for *bits*, *fields*, and *registers* that are important from a software perspective (for instance, address bits used by software, and programmable fields and registers), and various *floating point instruction formats*, such as *S* and *D*
- is used for the memory access types, such as *cached* and *uncached*

1.1.2 Bold Text

- represents a term that is being **defined**
- is used for **bits** and **fields** that are important from a hardware perspective (for instance, **register** bits, which are not programmable but accessible only to hardware)
- is used for ranges of numbers; the range is indicated by an ellipsis. For instance, **5..1** indicates numbers 5 through 1
- is used to emphasize **UNPREDICTABLE** and **UNDEFINED** behavior, as defined below.

1.1.3 Courier Text

`Courier` fixed-width font is used for text that is displayed on the screen, and for examples of code and instruction pseudocode.

1.2 UNPREDICTABLE and UNDEFINED

The terms **UNPREDICTABLE** and **UNDEFINED** are used throughout this book to describe the behavior of the processor in certain cases. **UNDEFINED** behavior or operations can occur only as the result of executing instructions in a privileged mode (i.e., in Kernel Mode or Debug Mode, or with the CP0 usable bit set in the Status register). Unprivileged software can never cause **UNDEFINED** behavior or operations. Conversely, both privileged and unprivileged software can cause **UNPREDICTABLE** results or operations.

1.2.1 UNPREDICTABLE

UNPREDICTABLE results may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. Software can never depend on results that are **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause a result to be generated or not. If a result is generated, it is **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause arbitrary exceptions.

UNPREDICTABLE results or operations have several implementation restrictions:

- Implementations of operations generating **UNPREDICTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode

- **UNPREDICTABLE** operations must not read, write, or modify the contents of memory or internal state which is inaccessible in the current processor mode. For example, **UNPREDICTABLE** operations executed in user mode must not access memory or internal state that is only accessible in Kernel Mode or Debug Mode or in another process
- **UNPREDICTABLE** operations must not halt or hang the processor

1.2.2 UNDEFINED

UNDEFINED operations or behavior may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. **UNDEFINED** operations or behavior may vary from nothing to creating an environment in which execution can no longer continue. **UNDEFINED** operations or behavior may cause data loss.

UNDEFINED operations or behavior has one implementation restriction:

- **UNDEFINED** operations or behavior must not cause the processor to hang (that is, enter a state from which there is no exit other than powering down the processor). The assertion of any of the reset signals must restore the processor to an operational state

1.2.3 UNSTABLE

UNSTABLE results or values may vary as a function of time on the same implementation or instruction. Unlike **UNPREDICTABLE** values, software may depend on the fact that a sampling of an **UNSTABLE** value results in a legal transient value that was correct at some point in time prior to the sampling.

UNSTABLE values have one implementation restriction:

- Implementations of operations generating **UNSTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode

1.3 Special Symbols in Pseudocode Notation

In this book, algorithmic descriptions of an operation are described using a high-level language pseudocode resembling Pascal. Special symbols used in the pseudocode notation are listed in Table 1.1.

Table 1.1 Symbols Used in Instruction Operation Statements

Symbol	Meaning
\leftarrow	Assignment
$=, \neq$	Tests for equality and inequality
\parallel	Bit string concatenation
x^y	A y -bit string formed by y copies of the single-bit value x
$b\#n$	A constant value n in base b . For instance $10\#100$ represents the decimal value 100, $2\#100$ represents the binary value 100 (decimal 4), and $16\#100$ represents the hexadecimal value 100 (decimal 256). If the "b#" prefix is omitted, the default base is 10.
$0bn$	A constant value n in base 2. For instance $0b100$ represents the binary value 100 (decimal 4).
$0xn$	A constant value n in base 16. For instance $0x100$ represents the hexadecimal value 100 (decimal 256).

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
$x_{y..z}$	Selection of bits y through z of bit string x . Little-endian bit notation (rightmost bit is 0) is used. If y is less than z , this expression is an empty (zero length) bit string.
$x.\text{bit}[y]$	Bit y of bitstring x . Alternative to the traditional MIPS notation x_y .
$x.\text{bits}[y..z]$	Selection of bits y through z of bit string x . Alternative to the traditional MIPS notation $x_{y..z}$.
$x.\text{byte}[y]$	Byte y of bitstring x . Equivalent to the traditional MIPS notation $x_{8*y+7..8*y}$.
$x.\text{bytes}[y..z]$	Selection of bytes y through z of bit string x . Alternative to the traditional MIPS notation $x_{8*y+7..8*z}$.
$x.\text{halfword}[y]$ $x.\text{word}[i]$ $x.\text{doubleword}[i]$	Similar extraction of particular bitfields (used in e.g., MSA packed SIMD vectors).
$x.\text{bit}31, x.\text{byte}0$, etc.	Examples of abbreviated form of $x.\text{bit}[y]$, etc. notation, when y is a constant.
$x.\text{fiel}dy$	Selection of a named subfield of bitstring x , typically a register or instruction encoding. More formally described as “Field y of register x ”. For example, FIR.D = “the D bit of the Coprocessor 1 Floating-point Implementation Register (FIR)”.
$+, -$	2’s complement or floating point arithmetic: addition, subtraction
$*, \infty$	2’s complement or floating point multiplication (both used for either)
div	2’s complement integer division
mod	2’s complement modulo
$/$	Floating point division
$<$	2’s complement less-than comparison
$>$	2’s complement greater-than comparison
\leq	2’s complement less-than or equal comparison
\geq	2’s complement greater-than or equal comparison
nor	Bitwise logical NOR
xor	Bitwise logical XOR
and	Bitwise logical AND
or	Bitwise logical OR
not	Bitwise inversion
$\&\&$	Logical (non-Bitwise) AND
\ll	Logical Shift left (shift in zeros at right-hand-side)
\gg	Logical Shift right (shift in zeros at left-hand-side)
$\text{GPRL}EN$	The length in bits (32 or 64) of the CPU general-purpose registers
$GPR[x]$	CPU general-purpose register x . The content of $GPR[0]$ is always zero. In Release 2 of the Architecture, $GPR[x]$ is a short-hand notation for $SGPR[SRSCtl_{CSS}, x]$.
$SGPR[s,x]$	In Release 2 of the Architecture and subsequent releases, multiple copies of the CPU general-purpose registers may be implemented. $SGPR[s,x]$ refers to GPR set s , register x .
$FPR[x]$	Floating Point operand register x
$FCC[CC]$	Floating Point condition code CC . $FCC[0]$ has the same value as $COC[1]$. Release 6 removes the floating point condition codes.
$FPR[x]$	Floating Point (Coprocessor unit 1), general register x

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
$CPR[z,x,s]$	Coprocessor unit z , general register x , select s
$CP2CPR[x]$	Coprocessor unit 2, general register x
$CCR[z,x]$	Coprocessor unit z , control register x
$CP2CCR[x]$	Coprocessor unit 2, control register x
$COC[z]$	Coprocessor unit z condition signal
$Xlat[x]$	Translation of the MIPS16e GPR number x into the corresponding 32-bit GPR number
BigEndianMem	Endian mode as configured at chip reset (0 → Little-Endian, 1 → Big-Endian). Specifies the endianness of the memory interface (see LoadMemory and StoreMemory pseudocode function descriptions) and the endianness of Kernel and Supervisor mode execution.
BigEndianCPU	The endianness for load and store instructions (0 → Little-Endian, 1 → Big-Endian). In User mode, this endianness may be switched by setting the <i>RE</i> bit in the <i>Status</i> register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).
ReverseEndian	Signal to reverse the endianness of load and store instructions. This feature is available in User mode only, and is implemented by setting the <i>RE</i> bit of the <i>Status</i> register. Thus, ReverseEndian may be computed as (SR _{RE} and User mode).
<i>LLbit</i>	Bit of virtual state used to specify operation for instructions that provide atomic read-modify-write. <i>LLbit</i> is set when a linked load occurs and is tested by the conditional store. It is cleared, during other CPU operation, when a store to the location would no longer be atomic. In particular, it is cleared by exception return instructions.
I , I+n , I-n :	This occurs as a prefix to <i>Operation</i> description lines and functions as a label. It indicates the instruction time during which the pseudocode appears to “execute.” Unless otherwise indicated, all effects of the current instruction appear to occur during the instruction time of the current instruction. No label is equivalent to a time label of I . Sometimes effects of an instruction appear to occur either earlier or later — that is, during the instruction time of another instruction. When this happens, the instruction operation is written in sections labeled with the instruction time, relative to the current instruction I , in which the effect of that pseudocode appears to occur. For example, an instruction may have a result that is not available until after the next instruction. Such an instruction has the portion of the instruction operation description that writes the result register in a section labeled I+1 . The effect of pseudocode statements for the current instruction labeled I+1 appears to occur “at the same time” as the effect of pseudocode statements labeled I for the following instruction. Within one pseudocode sequence, the effects of the statements take place in order. However, between sequences of statements for different instructions that occur “at the same time,” there is no defined order. Programs must not depend on a particular order of evaluation between such sections.
PC	The <i>Program Counter</i> value. During the instruction time of an instruction, this is the address of the instruction word. The address of the instruction that occurs during the next instruction time is determined by assigning a value to <i>PC</i> during an instruction time. If no value is assigned to <i>PC</i> during an instruction time by any pseudocode statement, it is automatically incremented by either 2 (in the case of a 16-bit MIPS16e instruction) or 4 before the next instruction time. A taken branch assigns the target address to the <i>PC</i> during the instruction time of the instruction in the branch delay slot. In the MIPS Architecture, the <i>PC</i> value is only visible indirectly, such as when the processor stores the restart address into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception. Release 6 adds PC-relative address computation and load instructions. The <i>PC</i> value contains a full 64-bit address, all of which are significant during a memory reference.

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning						
ISA Mode	<p>In processors that implement the MIPS16e Application Specific Extension or the microMIPS base architectures, the <i>ISA Mode</i> is a single-bit register that determines in which mode the processor is executing, as follows:</p> <table data-bbox="597 394 1266 541"> <tr> <th>Encoding</th><th>Meaning</th></tr> <tr> <td>0</td><td>The processor is executing 32-bit MIPS instructions</td></tr> <tr> <td>1</td><td>The processor is executing MIPS16e or microMIPS instructions</td></tr> </table> <p>In the MIPS Architecture, the <i>ISA Mode</i> value is only visible indirectly, such as when the processor stores a combined value of the upper bits of PC and the <i>ISA Mode</i> into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception.</p>	Encoding	Meaning	0	The processor is executing 32-bit MIPS instructions	1	The processor is executing MIPS16e or microMIPS instructions
Encoding	Meaning						
0	The processor is executing 32-bit MIPS instructions						
1	The processor is executing MIPS16e or microMIPS instructions						
PABITS	The number of physical address bits implemented is represented by the symbol PABITS. As such, if 36 physical address bits were implemented, the size of the physical address space would be $2^{\text{PABITS}} = 2^{36}$ bytes.						
SEGBITS	The number of virtual address bits implemented in a segment of the address space is represented by the symbol SEGBITS. As such, if 40 virtual address bits are implemented in a segment, the size of the segment is $2^{\text{SEGBITS}} = 2^{40}$ bytes.						
FP32RegistersMode	<p>Indicates whether the FPU has 32-bit or 64-bit floating point registers (FPRs). It is optional if the FPU has 32 64-bit FPRs in which 64-bit data types are stored in any FPR.</p> <p>microMIPS64 implementations have a compatibility mode in which the processor references the FPRs as if it were a microMIPS32 implementation. In such a case FP32RegisterMode is computed from the FR bit in the <i>Status</i> register. If this bit is a 0, the processor operates as if it had 32, 32-bit FPRs. If this bit is a 1, the processor operates with 32 64-bit FPRs.</p> <p>The value of FP32RegistersMode is computed from the FR bit in the <i>Status</i> register.</p>						
InstructionInBranchDelaySlot	Indicates whether the instruction at the Program Counter address was executed in the delay slot of a branch or jump. This condition reflects the <i>dynamic</i> state of the instruction, not the <i>static</i> state. That is, the value is false if a branch or jump occurs to an instruction whose PC immediately follows a branch or jump, but which is not executed in the delay slot of a branch or jump.						
SignalException(exception, argument)	Causes an exception to be signaled, using the exception parameter as the type of exception and the argument parameter as an exception-specific argument). Control does not return from this pseudocode function—the exception is signaled at the point of the call.						

1.4 Notation for Register Field Accessibility

In this document, the read/write properties of register fields use the notations shown in Table 1.1.

Table 1.2 Read/Write Register Field Notation

Read/Write Notation	Hardware Interpretation	Software Interpretation
R/W	<p>A field in which all bits are readable and writable by software and, potentially, by hardware. Hardware updates of this field are visible by software read. Software updates of this field are visible by hardware read.</p> <p>If the Reset State of this field is “Undefined”, either software or hardware must initialize the value before the first read will return a predictable value. This should not be confused with the formal definition of UNDEFINED behavior.</p>	

Table 1.2 Read/Write Register Field Notation (Continued)

Read/Write Notation	Hardware Interpretation	Software Interpretation
R	<p>A field which is either static or is updated only by hardware.</p> <p>If the Reset State of this field is either “0”, “Preset”, or “Externally Set”, hardware initializes this field to zero or to the appropriate state, respectively, on powerup. The term “Preset” is used to suggest that the processor establishes the appropriate state, whereas the term “Externally Set” is used to suggest that the state is established via an external source (e.g., personality pins or initialization bit stream). These terms are suggestions only, and are not intended to act as a requirement on the implementation.</p> <p>If the Reset State of this field is “Undefined”, hardware updates this field only under those conditions specified in the description of the field.</p>	<p>A field to which the value written by software is ignored by hardware. Software may write any value to this field without affecting hardware behavior. Software reads of this field return the last value updated by hardware.</p> <p>If the Reset State of this field is “Undefined”, software reads of this field result in an UNPREDICTABLE value except after a hardware update done under the conditions specified in the description of the field.</p>
R0	<p>R0 = reserved, read as zero, ignore writes by software.</p> <p>Hardware ignores software writes to an R0 field. Neither the occurrence of such writes, nor the values written, affects hardware behavior.</p> <p>Hardware always returns 0 to software reads of R0 fields.</p> <p>The Reset State of an R0 field must always be 0.</p> <p>If software performs an mtc0 instruction which writes a non-zero value to an R0 field, the write to the R0 field will be ignored, but permitted writes to other fields in the register will not be affected.</p>	<p>Architectural Compatibility: R0 fields are reserved, and may be used for not-yet-defined purposes in future revisions of the architecture.</p> <p>When writing an R0 field, current software should only write either all 0s, or, preferably, write back the same value that was read from the field.</p> <p>Current software should not assume that the value read from R0 fields is zero, because this may not be true on future hardware.</p> <p>Future revisions of the architecture may redefine an R0 field, but must do so in such a way that software which is unaware of the new definition and either writes zeros or writes back the value it has read from the field will continue to work correctly.</p> <p>Writing back the same value that was read is guaranteed to have no unexpected effects on current or future hardware behavior. (Except for non-atomicity of such read-writes.)</p> <p>Writing zeros to an R0 field may not be preferred because in the future this may interfere with the operation of other software which has been updated for the new field definition.</p>

Table 1.2 Read/Write Register Field Notation (Continued)

Read/Write Notation	Hardware Interpretation	Software Interpretation
0	Release 6 Release 6 legacy “0” behaves like R0 - read as zero, nonzero writes ignored. Legacy “0” should not be defined for any new control register fields; R0 should be used instead.	
	HW returns 0 when read. HW ignores writes.	Only zero should be written, or, value read from register.
	pre-Release 6 pre-Release 6 legacy “0” - read as zero, nonzero writes UNDEFINED	
	A field which hardware does not update, and for which hardware can assume a zero value.	A field to which the value written by software must be zero. Software writes of non-zero values to this field may result in UNDEFINED behavior of the hardware. Software reads of this field return zero as long as all previous software writes are zero. If the Reset State of this field is “Undefined”, software must write this field with zero before it is guaranteed to read as zero.
R/W0	Like R/W, except that writes of non-zero to a R/W0 field are ignored. E.g. Status.NMI	
	Hardware may set or clear an R/W0 bit. Hardware ignores software writes of nonzero to an R/W0 field. Neither the occurrence of such writes, nor the values written, affects hardware behavior. Software writes of 0 to an R/W0 field may have an effect. Hardware may return 0 or nonzero to software reads of an R/W0 bit. If software performs an mtc0 instruction which writes a non-zero value to an R/W0 field, the write to the R/W0 field will be ignored, but permitted writes to other fields in the register will not be affected.	Software can only clear an R/W0 bit. Software writes 0 to an R/W0 field to clear the field. Software writes nonzero to an R/W0 bit in order to guarantee that the bit is not affected by the write.

1.5 For More Information

MIPS processor manuals and additional information about MIPS products can be found at <http://www.imgtec.com>.

For comments or questions on the MIPS64® Architecture or this document, send Email to IMGBA-DocFeedback@imgtec.com.

Introduction

In today's market, the lowest price, performance, or both must be satisfied, especially for deeply-embedded applications such as microcontroller applications. Moreover, customers require efficient solutions that can be turned into products quickly. To meet this need, the MIPS® instruction set has been optimized and re-encoded into a new variable-length scheme. This solution is called microMIPS™.

microMIPS minimizes the resulting code footprint of applications and reduces the cost of memory, which is particularly high for embedded memory. Simultaneously, the high performance of MIPS cores is maintained. Using this technology, the customer can generate best results without spending time to profile its application. The smaller code footprint typically leads to reduced power consumption per executed task because of the smaller number of memory accesses.

microMIPS is the preferred replacement for the existing MIPS16e™ ASE. MIPS16e could only be used for user mode programs which did not use floating-point nor any of the Application Specific Extensions (ASEs). microMIPS does not have these limitations — it can be used for kernel mode code as well as user mode programs. It can be used for programs which use floating-point. It can be used with the available ASEs.

microMIPS is also an alternative to the MIPS64® instruction encoding and can be implemented in parallel or stand-alone. The microMIPS equivalent of MIPS32 is microMIPS32™ and the microMIPS equivalent of MIPS64 is microMIPS64™.

Overview of changes vs. existing MIPS64 ISA:

- 16-bit and 32-bit opcodes
- Optimized opcode/operand field definitions based on statistics
- Removal of branch likely instructions, emulation by assembler
- Fine-tuned register allocation algorithm in the compiler for lowest code size

2.1 Default ISA Mode

The instruction sets which are available within an implementation are reported by the *Config3_{ISA}* register field (bits 15:14). *Config1_{CA}* (bit 2) is not used for microMIPS64.

For implementations that support both microMIPS64 and MIPS64, the selected ISA mode following reset is determined by the setting of the *Config3_{ISA}* register field., which is a read-only field set by a hardware signal external to the processor core.

For implementations that support both microMIPS64 and MIPS64, the selected ISA mode upon handling an exception is determined by the setting of the *Config3_{ISAOnExc}* register field (bit 16). The *Config3_{ISAOnExc}* register field is writeable by software and has a reset value that is set by a hardware signal external to the processor core. This register

field allows privileged software to change the ISA mode to be used for subsequent exceptions. This capability is for all exception types whose vectors are offsets of the *EBASE* register.

For implementations that support both microMIPS64 and MIPS64, the selected ISA mode when handling a debug exception is determined by the setting of the *ISAonDebug* register field in the *EJTAG TAP Control* register. This register field is writeable by EJTAG probe software and has a reset value that is set by a hardware signal external to the processor core.

For CPU cores supporting the MT ASE and multiple VPEs, the ISA mode for exceptions can be selected on a per-VPE basis.

2.2 Software Detection

Software can determine if microMIPS64 ISA is implemented by checking the state of the ISA (Instruction Set Architecture) field in the *Config3* CP0 register. *Config1_{CA}* (bit 2) is not used for microMIPS64.

Software can determine if the MIPS64 ISA is implemented by checking the state of the ISA (Instruction Set Architecture) register field in the *Config3* CP0 register.

Software can determine which ISA is used when handling an exception by checking the state of the *ISAonExc* (ISA on Exception) field in the *Config3* CP0 register.

Debug Probe Software can determine which ISA is used when handling a debug exception by checking the state of the *ISAonDebug* field in the *EJTAG TAP Control* register.

2.3 Compliance and Subsetting

This document does not change the instruction subsets as defined by the other MIPS architecture reference manuals, including the subsets defined by the various ASEs.

2.4 ISA Mode Switch

The MIPS Release 3 architecture defines an ISA mode for each processor. An ISA mode value of 0 indicates MIPS64 instruction decoding. In processors implementing microMIPS64, an ISA mode value of 1 selects microMIPS64 instruction decoding.

The ISA mode is not directly visible to user mode software. Upon an exception, the ISA mode of the faulting/interrupted instruction is recorded in the least-significant address bit within the appropriate return address register - either *EPC* or *ErrorEPC* or *DebugEPC*, depending on the exception type.

For the rest of this section, the following definitions are used:

Jump-and-Link-Register instructions: For the MIPS64 ISA, this means the JALR and JALR.HB instructions. For the microMIPS64 ISA, this means the JALRC, JALRC.HB, JIALC, and JALRC16 instructions.

Jump-Register instructions: For the MIPS64 ISA, this means the JR and JR.HB instructions. For the microMIPS64 ISA, this means the instructions JRC, JRC.HB, JIC, JRC16, and JRCADDIUSP instructions.

Mode switching between MIPS64 and microMIPS64 is enabled by the Jump-and-Link-Register and Jump-Register instructions, as described below.

- The Jump-and-Link-Register and Jump-Register instructions interpret bit 0 of the source registers as the target ISA mode (0=MIPS64, 1=microMIPS64) and therefore set the ISA Mode bit according to the contents of bit 0 of the source register. For the actual jump operation, the PC is loaded with the value of the source register with bit 0 set to 0. The Jump-and-Link-Register instructions save the ISA mode into bit 0 of the destination register.
- When exceptions or interrupts occur and the processor writes to *EPC*, *DEPC*, or *ErrorEPC*, the ISA Mode bit is saved into bit 0 of these registers. Then the ISA Mode bit is set according to the *Config3_{ISAOnExc}* register field. On return from an exception, the processor loads the ISA Mode bit based on the value from either *EPC*, *DEPC*, or *ErrorEPC*.

If only one ISA mode exists (either MIPS64 or microMIPS64) then this mode switch mechanism does not exist, but the ISA Mode bit is still maintained and has a fixed value (0=MIPS64, 1=microMIPS64). This is to maintain code compatibility between devices which implement both ISA modes and devices which implement only one ISA mode. Jump-Register and Jump-and-Link-Register instructions cause an Address exception on the target instruction fetch when bit 0 of the source register is different from the fixed ISA mode. Exception handlers must use the instruction set binary format supported by the processor. The Jump-and-Link-Register instructions must still save the fixed ISA mode into bit 0 of the destination register.

2.5 Branch and Jump Offsets

In the MIPS64 architecture, because instructions are always 32 bits in size, the jump and branch target addresses are word (32-bit) aligned. Jump/branch offset fields are shifted left by two bits to create a word-aligned effective address.

In the microMIPS64 architecture, because instructions can be either 16 or 32 bits in size, the jump and branch target addresses are halfword (16-bit) aligned. Branch/jump offset fields are shifted left by only one bit to create halfword-aligned effective addresses.

To maintain the existing MIPS64 ABIs, link unit/object file entry points are restricted to 32-bit word alignments. In the future, a microMIPS64-only ABI can be created to remove this restriction.

2.6 Coprocessor Unusable Behavior

If an instruction associated with a non-implemented coprocessor is executed, it is implementation specific whether a processor executing in microMIPS64 mode raises an RI exception or a coprocessor unusable exception. This behavior is different from the MIPS64 behavior in which coprocessor unusable exception is signalled for such cases.

If the microMIPS64 implementation chooses to use RI exception in such cases, the microMIPS64 RI exception handler must check for coprocessor instructions being executed while the associated coprocessor is implemented but has been disabled (*Status_{CUx}* set to zero).

2.7 Release 6 of the MIPS Architecture

Table 2.1 Instructions Added in Release 6

Instruction	Instruction's Purpose	Replaces
ADDIUPC	Add Immediate to PC (unsigned - non-trapping)	New
ALIGN	Concatenate two GPRs, and extract a contiguous subset at a byte position (32-bit)	New
DALIGN	Concatenate two GPRs, and extract a contiguous subset at a byte position (64-bit)	New
ALUIPC	Aligned Add Upper Immediate to PC	New
AUI	Add Upper Immediate	New
DAUI	Doubleword Add Upper Immediate	New
DAHI	Doubleword Add Higher Immediate	New
DATI	Doubleword Add Top Immediate	New
AUIPC	Add Upper Immediate to PC	New
BC1EQZC	Branch if Coprocessor 1 (FPU) Register Bit 0 is Equal to Zero	BC1F
BC1NEZC	Branch if Coprocessor 1 (FPR) Register Bit 0 is Not Equal to Zero	BC1T
BC2EQZC	Branch if Coprocessor 2 Condition Register is Equal to Zero	BC2F
BC2NEZC	Branch if Coprocessor 2 Condition Register is Not Equal to Zero	BC2T
BLEZALC	Compact branch-and-link if GPR rt is less than or equal to zero	New
BGEZALC	Compact branch-and-link if GPR rt is greater than or equal to zero	Compact version
BGTZALC	Compact branch-and-link if GPR rt is greater than zero	New
BLTZALC	Compact branch-and-link if GPR rt is less than to zero	Compact version
BEQZALC	Compact branch-and-link if GPR rt is equal to zero	New
BNEZALC	Compact branch-and-link if GPR rt is not equal to zero	New
BEQC	Equal register-register compare and branch with 16-bit offset	New
BNEC	Not-Equal register-register compare and branch with 16-bit offset	New
BLTC	Signed register-register compare and branch with 16-bit offset:67	New
BGEC	Signed register-register compare and branch with 16-bit offset:	New
BLTUC	Unsigned register-register compare and branch with 16-bit offset:	New
BGEUC	Unsigned register-register compare and branch with 16-bit offset:	New
BGTC	Assembly idioms with reversed operands for signed/unsigned compare-and-branch	New
BLEC	Assembly idioms with reversed operands for signed/unsigned compare-and-branch	New
BGTUC	Assembly idioms with reversed operands for signed/unsigned compare-and-branch	New
BLEUC	Assembly idioms with reversed operands for signed/unsigned compare-and-branch	New
BLTZC	Signed Compare register to Zero and branch with 16-bit offset	Compact version

Table 2.1 Instructions Added in Release 6 (Continued)

Instruction	Instruction's Purpose	Replaces
BLEZC	Signed Compare register to Zero and branch with 16-bit offset	Compact version
BGEZC	Signed Compare register to Zero and branch with 16-bit offset	Compact version
BGTZC	Signed Compare register to Zero and branch with 16-bit offset	Compact version
BEQZC	Equal Compare register to Zero and branch with 21-bit offset	Compact version with 21-bit offset
BNEZC	Not-equal Compare register to Zero and branch with 21-bit offset	Compact version with 21-bit offset
BC/BC16	Branch, Compact (16)	B/B16
BALC	Branch and Link, Compact	BAL
BITSWAP	Swaps (reverses) bits in each byte	New
DBITSWAP	Swaps (reverses) bits in each byte	New
BOVC	Branch on Overflow, Compact; Branch on No Overflow, Compact	New
BNVC	Branch on Overflow, Compact; Branch on No Overflow, Compact	New
CRC32B/CRC32H/ CRC32W/CRC32D	Generate CRC with reversed polynomial 0xEDB88320	New
CRC32CB/CRC32CH/ CRC32CW/CRC32CD	Generate CRC with reversed polynomial 0x82F63B78	New
CLASS.fmt	Scalar Floating-Point Class Mask	New
CMP.condn.fmt	Floating Point Compare setting Mask	C.condn.fmt
DIV	Divide Words Signed	DIV
DVP	Disable Virtual Processor	New
EVP	Enable Virtual Processor	New
MOD	Modulo Words Signed	DIV
DIVU	Divide Words Signed	DIVU
MODU	Modulo Words Signed	DIVU
DDIV	Divide Doublewords Signed	DDIV
DMOD	Modulo Doublewords Signed	DDIV
DDIVU	Divide Doublewords Signed	DDIVU
DMODU	Modulo Doublewords Signed	DDIVU
GINVI	Global Invalidate Instruction Cache	New
GINVT	Global Invalidate TLB	New
JALRC16	Jump and Link Register Compact (16-bit instr size)	JALR16
JIALC	Jump Indexed and Link, Compact	New
JIC	Jump Indexed, Compact	New
JRCADDIUSP	Jump Register, Adjust Stack Pointer (16-bit)	JRADDIUSP
LDPC	Load Doubleword PC-relative	New
LSA	Load Scaled Address	New
DLSA	Doubleword Load Scaled Address	New
LWPC	Load Word PC-relative	New
LWUPC	Load Word Unsigned PC-relative	New

Table 2.1 Instructions Added in Release 6 (Continued)

Instruction	Instruction's Purpose	Replaces
MADDF.fmt	Floating Point Fused Multiply Add	MADD.fmt
MSUBF.fmt	Floating Point Fused Multiply Subtract	MSUB.fmt
MAX.fmt	Scalar Floating-Point Maximum	New
MAXA.fmt	Scalar Floating-Point Argument with Maximum Absolute Value	New
MIN.fmt	Scalar Floating-Point Minimum	New
MINA.fmt	Scalar Floating-Point Argument with Minimum Absolute Value	New
MUL	Multiply Words Signed, Low Word	MULT
MUH	Multiply Words Signed, High Word	MULT
MULU	Multiply Words Signed, Low Word	MULTU
MUHU	Multiply Words Signed, High Word	MULTU
DMUL	Multiply Doublewords Signed, Low Doubleword	DMULT
DMUH	Multiply Doublewords Signed, High Doubleword	DMULT
DMULU	Multiply Doublewords Signed, Low Doubleword	DMULTU
DMUHU	Multiply Doublewords Signed, High Doubleword	DMULTU
RINT.fmt	Floating-Point Round to Integral	New
SEL.fmt	Select floating point values with FPR condition	MOVE.fmt, MOVT.fmt
SELEQZ	Select integer GPR value or zero	MOVZ, MOVN
SELNEZ	Select integer GPR value or zero	MOVZ, MOVN
SELEQZ.fmt	Select floating point value or zero with FPR condition	MOVZ.fmt, MOVN.fmt
SELNEZ.fmt	Select floating point value or zero with FPR condition	MOVZ.fmt, MOVN.fmt

Table 2.2 Instructions Recoded in Release 6

Instruction	Purpose
AND16	To do a bitwise logical AND
BEQZC	Branch on Equal to Zero, Compact
BNEZC	Branch on Not Equal to Zero, Compact
BREAK16	Breakpoint
JRC16	Jump Register, Compact (16-bit)
LUI	To load a constant into the upper half of a word
LWM16	Load Word Multiple (16-bit)
MOVEP	Move a Pair of Registers
NOT16	Invert (16-bit instr size)
OR16	Or (16-bit instr size)
SDBBP16	Software Debug Breakpoint (16-bit instr size)

Table 2.2 Instructions Recoded in Release 6 (Continued)

Instruction	Purpose
SWM16	Store Word Multiple (16-bit)
SYNCI	Synchronize Caches to Make Instruction Writes Effective
XOR16	Exclusive OR (16-bit instr size)

Table 2.3 Instructions Removed in Release 6

Instruction	Purpose	Replaced by
ABS.PS	Floating Point Absolute Value, Paired Single	—
ADD.PS	Floating Point Add, Paired Single	—
ADDI	Add Immediate Word	—
ALNV.PS	Floating Point Align Variable, Paired Single	—
B	Unconditional Branch	BC
B16	Unconditional Branch (16-bit instr size)	BC16
BAL	Branch and Link	BALC
BC1F	Branch on FP False	BC1EQZC
BC1T	Branch on FP True	BC1NEZC
BC2F	Branch on COP2 False	BC2EQZC
BC2T	Branch on COP2 True	BC2NEZC
BEQ	Branch on Equal	BEQC
BGEZ	Branch on Greater Than or Equal to Zero	BGEZC
BEQZ16	Branch on Equal to Zero (16-bit instr size)	BEQZC16
BGEZAL	Branch on Greater Than or Equal to Zero and Link	BGEZALC
BGEZALS	Branch on Greater Than or Equal to Zero and Link, Short Delay-Slot	—
BGTZ	Branch on Greater Than Zero	BGTZC
BLEZ	Branch on Less Than or Equal to Zero	BLEZC
BLTZ	Branch on Less Than Zero	BLTZC
BLTZAL	Branch on Less Than Zero and Link	BLTZALC
BLTZALS	Branch on Less Than Zero and Link, Short Delay-Slot	BLTZALC.
BNE	Branch on Not Equal	BNEC
BNEZ16	Branch on Not Equal to Zero (16-bit instr size)	BNEZC16.
C.cond.fmt	Floating Point Compare	CMP.condn.fmt

Table 2.3 Instructions Removed in Release 6 (Continued)

Instruction	Purpose	Replaced by
CVT.PS.S	Floating Point Convert Pair to Paired Single	—
CVT.S.PL	Floating Point Convert Pair Lower to Single Floating Point	—
CVT.S.PU	Floating Point Convert Pair Upper to Single Floating Point	—
DADDI	Doubleword Add Immediate	—
DDIV	Doubleword Divide	—
DDIVU	Doubleword Divide Unsigned	—
DIV	Divide Word	—
DIVU	Divide Unsigned Word	—
DMULT	Doubleword Multiply	—
DMULTU	Doubleword Multiply Unsigned	—
JALC	Jump and Link Compact	—
JALR16	Jump and Link Register (16-bit instr size)	JALRC16
JALRS	Jump and Link Register, Short Delay Slot	JALRC
JALRS.HB	Jump and Link Register with Hazard Barrier, Short Delay-Slot	—
JALRS16	Jump and Link Register, Short Delay-Slot (16-bit instr size)	JALRC16
JALS	Jump and Link, Short Delay Slot	—
JALX	Jump and Link Exchange (microMIPS Format)	—
JC	Jump Register, Compact	—
JR	Jump Register	JALRC
JR.HB	Jump Register with Hazard Barrier	JALRC.HB
JRC	Jump Register, Compact (16)	—
JR16	Jump Register (16-bit instr size)	JRC16
JRADDIUSP	Jump Register, Adjust Stack Pointer	JRCADDIUSP
LDL	Load Doubleword Left	—
LDR	Load Doubleword Right	—
LDXC1	Load Doubleword Indexed to Floating Point	—
LUXC1	Load Doubleword Indexed Unaligned to Floating Point	—
LWL	Load Word Left	—
LWLE	Load Word Left EVA	—
LWR	Load Word Right	—
LWRE	Load Word Right EVA	—

Table 2.3 Instructions Removed in Release 6 (Continued)

Instruction	Purpose	Replaced by
LWXC1	Load Word Indexed to Floating Point	—
LWXS	Load Word Indexed, Scaled	—
MADD	Multiply and Add Word to Hi, Lo	—
MADD.fmt	Floating Point Multiply Add	MADDF.fmt
MADDU	Multiply and Add Unsigned Word to Hi,Lo	—
MFHI16	Move From HI Register (16-bit instr size)	—
MFLO16	Move From LO Register	—
MFHI	Move From HI Register	—
MFLO	Move From LO Register	—
MOV.PS	Floating Point Move	—
MOV.F.fmt	Floating Point Move Conditional on Floating Point False	SEL.fmt
MOVN	Move Conditional on Not Zero	SELNEZ, SELEQZ
MOVN.fmt	Floating Point Move Conditional on Not Zero	SELNEZ.fmt
MOVT	Move Conditional on Floating Point True	—
MOVT.fmt	Floating Point Move Conditional on Floating Point True	SEL.fmt
MOVZ	Move Conditional on Zero	SELNEZ, SELEQZ
MOVZ.fmt	Floating Point Move Conditional on Zero	SELEZQZ.fmt
MSUB	Multiply and Subtract Word to Hi, Lo	—
MSUB.fmt	Floating Point Multiply Subtract	MSUBF.fmt
MSUBU	Multiply and Subtract Word to Hi,Lo	—
MTHI	Move to HI Register	—
MTLO	Move to LO Register	—
MUL	Multiply Word to GPR	—
MUL.PS	Floating Point Multiply, Paired Single	—
MULT	Multiply Word	MUL, MULH
MULTU	Multiply Unsigned Word	MULU, MUHU
NEG.PS	Floating Point Negate, Paired Single	—
NMADD.fmt	Floating Point Negative Multiply Add	NMADDF.fmt
NMSUB.fmt	Floating Point Negative Multiply Subtract	NMSUBF.fmt
PLL.PS	Pair Lower Lower, Paired Single	—
PLU.PS	Pair Lower Upper, Paired Single	—

Table 2.3 Instructions Removed in Release 6 (Continued)

Instruction	Purpose	Replaced by
PREFX	Prefetch Indexed	—
PUL.PS	Pair Upper Lower, Paired Single	—
PUU.PS	Pair Upper Upper, Paired Single	—
SDL	Store Doubleword Left	—
SDR	Store Doubleword Right	—
SDXC1	Store Doubleword Indexed from Floating Point	—
SUB.PS	Floating Point Subtract	—
SUXC1	Store Doubleword Indexed Unaligned from Floating Point	—
SWL	Store Word Left	—
SWLE	Store Word Left EVA	—
SWR	Store Word Right	—
SWXC1	Store Word Indexed from Floating Point	—
TEQI	Trap if Equal Immediate	—
TGEI	Trap if Greater or Equal Immediate	—
TGEIU	Trap if Greater or Equal Immediate Unsigned	—
TLTI	Trap if Less Than Immediate	—
TLTIU	Trap if Less Than Immediate Unsigned	—
TNEI	Trap if Not Equal Immediate	—

Guide to the Instruction Set

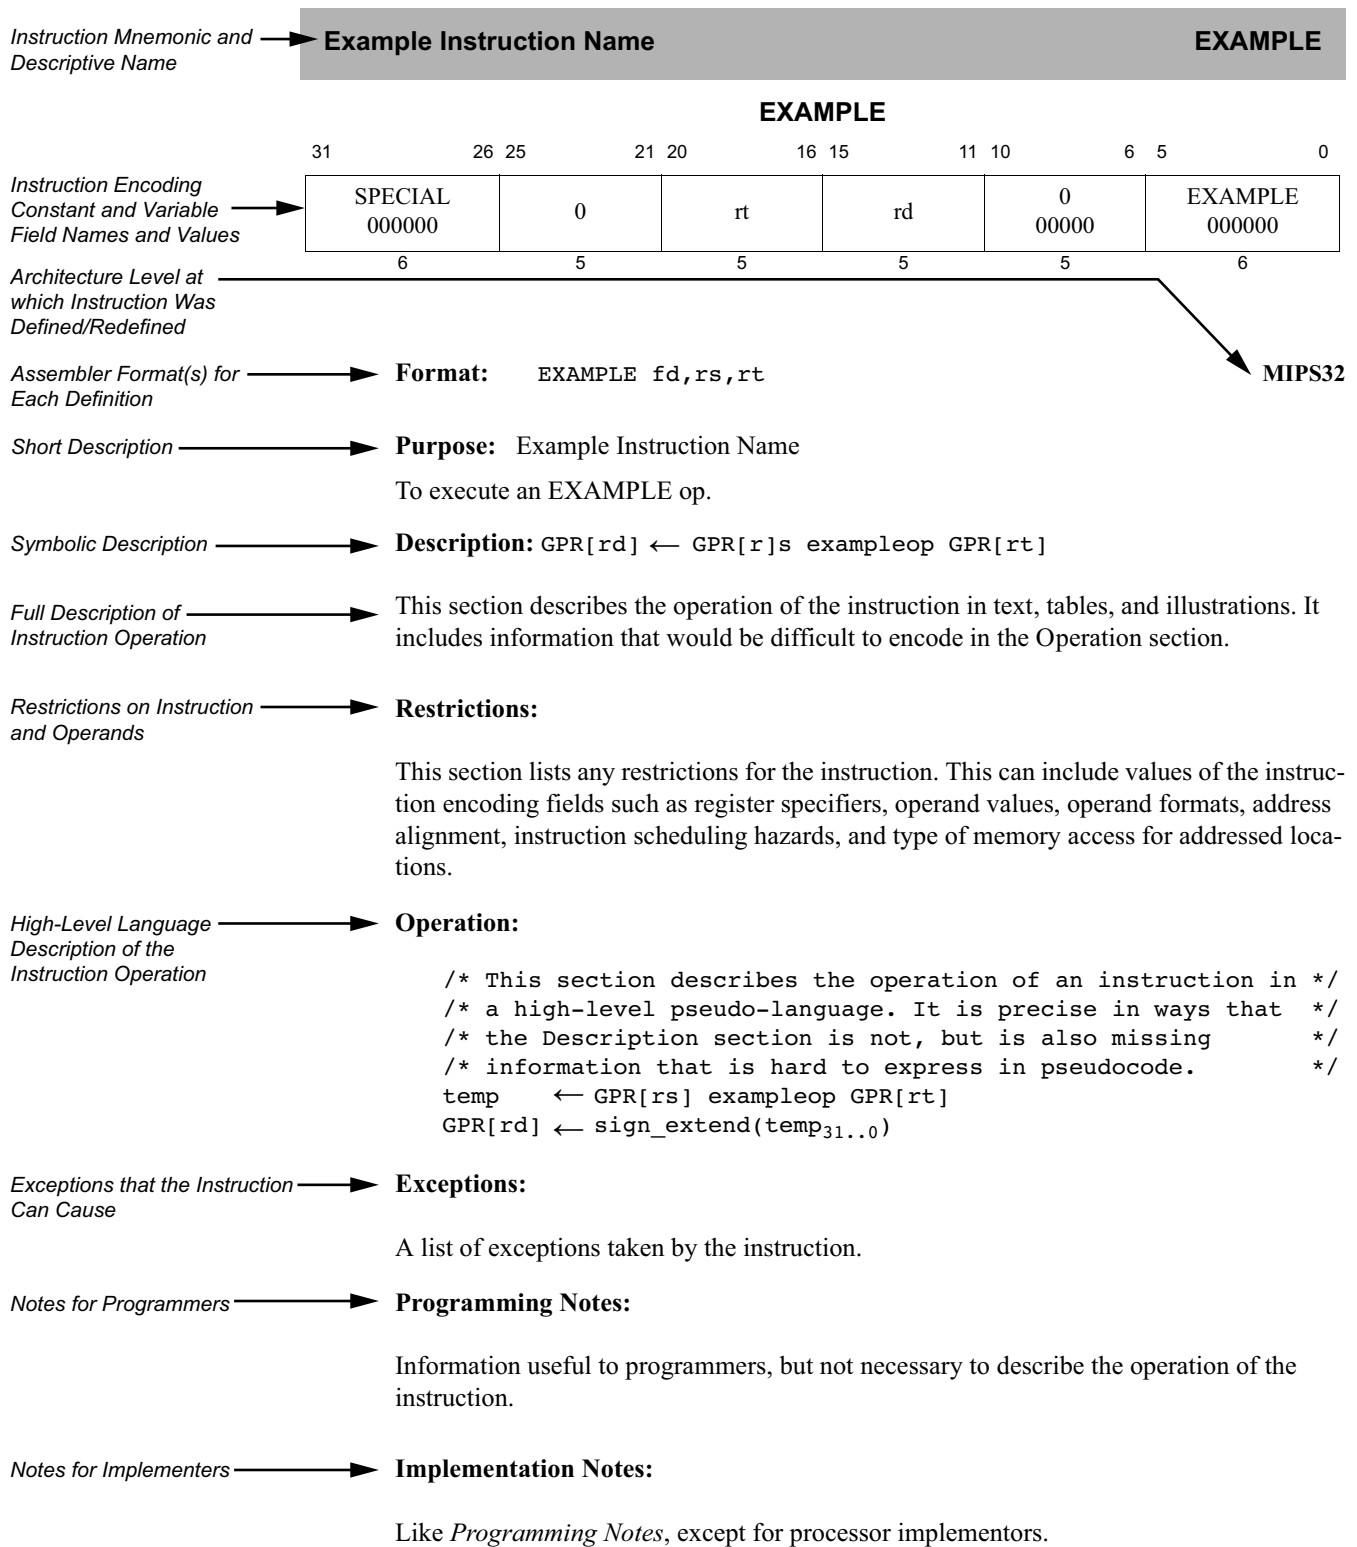
This chapter provides a detailed guide to understanding the instruction descriptions, which are listed in alphabetical order in the tables at the beginning of the next chapter.

3.1 Understanding the Instruction Fields

Figure 3.1 shows an example instruction. Following the figure are descriptions of the fields listed below:

- “Instruction Fields” on page 35
- “Instruction Descriptive Name and Mnemonic” on page 35
- “Format Field” on page 35
- “Purpose Field” on page 36
- “Description Field” on page 36
- “Restrictions Field” on page 37
- “Operation Field” on page 38
- “Exceptions Field” on page 38
- “Programming Notes and Implementation Notes Fields” on page 38

Figure 3.1 Example of Instruction Description



3.1.1 Instruction Fields

Fields encoding the instruction word are shown in register form at the top of the instruction description. The following rules are followed:

- The values of constant fields and the *opcode* names are listed in uppercase (SPECIAL and ADD in Figure 3.2). Constant values in a field are shown in binary below the symbolic or hexadecimal value.
- All variable fields are listed with the lowercase names used in the instruction description (*rs*, *rt*, and *rd* in Figure 3.2).
- Fields that contain zeros but are not named are unused fields that are required to be zero (bits 10:6 in Figure 3.2). If such fields are set to non-zero values, the operation of the processor is **UNPREDICTABLE**.

Figure 3.2 Example of Instruction Fields

31	26	25	21	20	16	15	11	10	6	5	0
SPECIAL 000000						rs		rt		rd	
0						00000		ADD		100000	
6						5		5		5	

3.1.2 Instruction Descriptive Name and Mnemonic

The instruction descriptive name and mnemonic are printed as page headings for each instruction, as shown in Figure 3.3.

Figure 3.3 Example of Instruction Descriptive Name and Mnemonic

Add Word	ADD
----------	-----

3.1.3 Format Field

The assembler formats for the instruction and the architecture level at which the instruction was originally defined are given in the *Format* field. If the instruction definition was later extended, the architecture levels at which it was extended and the assembler formats for the extended definition are shown in their order of extension (for an example, see C.cond.fmt). The MIPS architecture levels are inclusive; higher architecture levels include all instructions in previous levels. Extensions to instructions are backwards compatible. The original assembler formats are valid for the extended architecture.

Figure 3.4 Example of Instruction Format

Format:	ADD fd,rs,rt	MIPS32
----------------	--------------	---------------

The assembler format is shown with literal parts of the assembler instruction printed in uppercase characters. The variable parts, the operands, are shown as the lowercase names of the appropriate fields.

The architectural level at which the instruction was first defined, for example “MIPS32” is shown at the right side of the page. Instructions introduced at different times by different ISA family members, are indicated by markings such

as “MIPS64, MIPS32 Release 2”. Instructions removed by particular architecture release are indicated in the Availability section.

There can be more than one assembler format for each architecture level. Floating point operations on formatted data show an assembly format with the actual assembler mnemonic for each valid value of the *fmt* field. For example, the ADD.fmt instruction lists both ADD.S and ADD.D.

The assembler format lines sometimes include parenthetical comments to help explain variations in the formats (once again, see C.cond.fmt). These comments are not a part of the assembler format.

The term *decoded_immediate* is used if the immediate field is encoded within the binary format but the assembler format uses the decoded value. The term *left_shifted_offset* is used if the offset field is encoded within the binary format but the assembler format uses value after the appropriate amount of left shifting.

3.1.4 Purpose Field

The *Purpose* field gives a short description of the use of the instruction.

Figure 3.5 Example of Instruction Purpose

Purpose: Add Word

To add 32-bit integers. If an overflow occurs, then trap.

3.1.5 Description Field

If a one-line symbolic description of the instruction is feasible, it appears immediately to the right of the *Description* heading. The main purpose is to show how fields in the instruction are used in the arithmetic or logical operation.

Figure 3.6 Example of Instruction Description

Description: $\text{GPR}[rd] \leftarrow \text{GPR}[rs] + \text{GPR}[rt]$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs* to produce a 32-bit result.

- If the addition results in 32-bit 2’s complement arithmetic overflow, the destination register is not modified and an Integer Overflow exception occurs.
- If the addition does not overflow, the 32-bit result is signed-extended and placed into GPR *rd*.

The body of the section is a description of the operation of the instruction in text, tables, and figures. This description complements the high-level language description in the *Operation* section.

This section uses acronyms for register descriptions. “GPR *rt*” is CPU general-purpose register specified by the instruction field *rt*. “FPR *fs*” is the floating point operand register specified by the instruction field *fs*. “CP1 register *fd*” is the coprocessor 1 general register specified by the instruction field *fd*. “FCSR” is the floating point *Control / Status* register.

3.1.6 Restrictions Field

The *Restrictions* field documents any possible restrictions that may affect the instruction. Most restrictions fall into one of the following six categories:

- Valid values for instruction fields (for example, see floating point ADD.fmt)
- ALIGNMENT requirements for memory addresses (for example, see LW)
- Valid values of operands (for example, see ALNV.PS)
- Valid operand formats (for example, see floating point ADD.fmt)
- Order of instructions necessary to guarantee correct execution. These ordering constraints avoid pipeline hazards for which some processors do not have hardware interlocks (for example, see MUL).
- Valid memory access types (for example, see LL/SC)

Figure 3.7 Example of Instruction Restrictions

Restrictions:

If either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits _{63..31} equal), then the result of the operation is UNPREDICTABLE.

3.1.7 Availability and Compatibility Fields

The *Availability* and *Compatibility* sections are not provided for all instructions. These sections list considerations relevant to whether and how an implementation may implement some instructions, when software may use such instructions, and how software can determine if an instruction or feature is present. Such considerations include:

- Some instructions are not present on all architecture releases. Sometimes the implementation is required to signal a Reserved Instruction exception, but sometimes executing such an instruction encoding is architecturally defined to give UNPREDICTABLE results.
- Some instructions are available for implementations of a particular architecture release, but may be provided only if an optional feature is implemented. Control register bits typically allow software to determine if the feature is present.
- Some instructions may not behave the same way on all implementations. Typically this involves behavior that was UNPREDICTABLE in some implementations, but which is made architectural and guaranteed consistent so that software can rely on it in subsequent architecture releases.
- Some instructions are prohibited for certain architecture releases and/or optional feature combinations.
- Some instructions may be removed for certain architecture releases. Implementations may then be required to signal a Reserved Instruction exception for the removed instruction encoding; but sometimes the instruction encoding is reused for other instructions.

All of these considerations may apply to the same instruction. If such considerations applicable to an instruction are simple, the architecture level in which an instruction was defined or redefined in the *Format* field, and/or the *Restrictions* section, may be sufficient; but if the set of such considerations applicable to an instruction is complicated, the *Availability* and *Compatibility* sections may be provided.

3.1.8 Operation Field

The *Operation* field describes the operation of the instruction as pseudocode in a high-level language notation resembling Pascal. This formal description complements the *Description* section; it is not complete in itself because many of the restrictions are either difficult to include in the pseudocode or are omitted for legibility.

Figure 3.8 Example of Instruction Operation

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← (GPR[rs]31 | GPR[rs]31..0) + (GPR[rt]31 | GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← sign_extend(temp31..0)
endif

```

See 3.2 “Operation Section Notation and Functions” on page 39 for more information on the formal notation used here.

3.1.9 Exceptions Field

The *Exceptions* field lists the exceptions that can be caused by *Operation* of the instruction. It omits exceptions that can be caused by the instruction fetch, for instance, TLB Refill, and also omits exceptions that can be caused by asynchronous external events such as an Interrupt. Although a Bus Error exception may be caused by the operation of a load or store instruction, this section does not list Bus Error for load and store instructions because the relationship between load and store instructions and external error indications, like Bus Error, are dependent upon the implementation.

Figure 3.9 Example of Instruction Exception

Exceptions:

Integer Overflow

An instruction may cause implementation-dependent exceptions that are not present in the *Exceptions* section.

3.1.10 Programming Notes and Implementation Notes Fields

The *Notes* sections contain material that is useful for programmers and implementors, respectively, but that is not necessary to describe the instruction and does not belong in the description sections.

Figure 3.10 Example of Instruction Programming Notes**Programming Notes:**

ADDU performs the same arithmetic operation but does not trap on overflow.

3.2 Operation Section Notation and Functions

In an instruction description, the *Operation* section uses a high-level language notation to describe the operation performed by each instruction. Special symbols used in the pseudocode are described in the previous chapter. Specific pseudocode functions are described below.

This section presents information about the following topics:

- “Instruction Execution Ordering” on page 39
- “Pseudocode Functions” on page 39

3.2.1 Instruction Execution Ordering

Each of the high-level language statements in the *Operations* section are executed sequentially (except as constrained by conditional and loop constructs).

3.2.2 Pseudocode Functions

There are several functions used in the pseudocode descriptions. These are used either to make the pseudocode more readable, to abstract implementation-specific behavior, or both. These functions are defined in this section, and include the following:

- “Coprocessor General Register Access Functions” on page 39
- “Memory Operation Functions” on page 41
- “Floating Point Functions” on page 44
- “Miscellaneous Functions” on page 49

3.2.2.1 Coprocessor General Register Access Functions

Defined coprocessors, except for CP0, have instructions to exchange words and doublewords between coprocessor general registers and the rest of the system. What a coprocessor does with a word or doubleword supplied to it and how a coprocessor supplies a word or doubleword is defined by the coprocessor itself. This behavior is abstracted into the functions described in this section.

3.2.2.1.1 COP_LW

The COP_LW function defines the action taken by coprocessor *z* when supplied with a word from memory during a load word operation. The action is coprocessor-specific. The typical action would be to store the contents of memword in coprocessor general register *rt*.

Figure 3.11 COP_LW Pseudocode Function

```
COP_LW (z, rt, memword)
```

```

z: The coprocessor unit number
rt: Coprocessor general register specifier
memword: A 32-bit word value supplied to the coprocessor

/* Coprocessor-dependent action */

endfunction COP_LW

```

3.2.2.1.2 COP_LD

The COP_LD function defines the action taken by coprocessor *z* when supplied with a doubleword from memory during a load doubleword operation. The action is coprocessor-specific. The typical action would be to store the contents of memdouble in coprocessor general register *rt*.

Figure 3.12 COP_LD Pseudocode Function

```

COP_LD (z, rt, memdouble)
z: The coprocessor unit number
rt: Coprocessor general register specifier
memdouble: 64-bit doubleword value supplied to the coprocessor.

/* Coprocessor-dependent action */

endfunction COP_LD

```

3.2.2.1.3 COP_SW

The COP_SW function defines the action taken by coprocessor *z* to supply a word of data during a store word operation. The action is coprocessor-specific. The typical action would be to supply the contents of the low-order word in coprocessor general register *rt*.

Figure 3.13 COP_SW Pseudocode Function

```

dataword ← COP_SW (z, rt)
z: The coprocessor unit number
rt: Coprocessor general register specifier
dataword: 32-bit word value

/* Coprocessor-dependent action */

endfunction COP_SW

```

3.2.2.1.4 COP_SD

The COP_SD function defines the action taken by coprocessor *z* to supply a doubleword of data during a store doubleword operation. The action is coprocessor-specific. The typical action would be to supply the contents of the low-order doubleword in coprocessor general register *rt*.

Figure 3.14 COP_SD Pseudocode Function

```

datadouble ← COP_SD (z, rt)
z: The coprocessor unit number
rt: Coprocessor general register specifier
datadouble: 64-bit doubleword value

/* Coprocessor-dependent action */

```



```
endfunction COP_SD
```

3.2.2.1.5 CoprocessorOperation

The CoprocessorOperation function performs the specified Coprocessor operation.

Figure 3.15 CoprocessorOperation Pseudocode Function

```
CoprocessorOperation (z, cop_fun)

/* z:          Coprocessor unit number */
/* cop_fun:    Coprocessor function from function field of instruction */

/* Transmit the cop_fun value to coprocessor z */

endfunction CoprocessorOperation
```

3.2.2.2 Memory Operation Functions

Regardless of byte ordering (big- or little-endian), the address of a halfword, word, or doubleword is the smallest byte address of the bytes that form the object. For big-endian ordering this is the most-significant byte; for a little-endian ordering this is the least-significant byte.

In the *Operation* pseudocode for load and store operations, the following functions summarize the handling of virtual addresses and the access of physical memory. The size of the data item to be loaded or stored is passed in the *AccessLength* field. The valid constant names and values are shown in Table 3.1. The bytes within the addressed unit of memory (word for 32-bit processors or doubleword for 64-bit processors) that are used can be determined directly from the *AccessLength* and the two or three low-order bits of the address.

3.2.2.2.1 Misaligned Support

MIPS processors originally required all memory accesses to be naturally aligned. MSA (the MIPS SIMD Architecture) supported misaligned memory accesses for its 128 bit packed SIMD vector loads and stores, from its introduction in MIPS Release 5. Release 6 requires systems to provide support for misaligned memory accesses for all ordinary memory reference instructions: the system must provide a mechanism to complete a misaligned memory reference for this instruction, ranging from full execution in hardware to trap-and-emulate.

The pseudocode function MisalignedSupport encapsulates the version number check to determine if misalignment is supported for an ordinary memory access.

Figure 3.16 MisalignedSupport Pseudocode Function

```
predicate ← MisalignedSupport ()
return Config.AR ≥ 2 // Architecture Revision 2 corresponds to MIPS Release 6.
end function
```

See Appendix B, “Misaligned Memory Accesses” on page 511 for a more detailed discussion of misalignment, including pseudocode functions for the actual misaligned memory access.

3.2.2.2.2 AddressTranslation

The AddressTranslation function translates a virtual address to a physical address and its cacheability and coherency attribute, describing the mechanism used to resolve the memory reference.

Given the virtual address *vAddr*, and whether the reference is to Instructions or Data (*IorD*), find the corresponding physical address (*pAddr*) and the cacheability and coherency attribute (*CCA*) used to resolve the reference. If the virtual address is in one of the unmapped address spaces, the physical address and *CCA* are determined directly by the virtual address. If the virtual address is in one of the mapped address spaces then the TLB or fixed mapping MMU determines the physical address and access type; if the required translation is not present in the TLB or the desired access is not permitted, the function fails and an exception is taken.

Figure 3.17 AddressTranslation Pseudocode Function

```
(pAddr, CCA) ← AddressTranslation (vAddr, IorD, LorS)

/* pAddr: physical address */
/* CCA: Cacheability&Coherency Attribute, the method used to access caches */
/*      and memory and resolve the reference */

/* vAddr: virtual address */
/* IorD: Indicates whether access is for INSTRUCTION or DATA */
/* LorS: Indicates whether access is for LOAD or STORE */

/* See the address translation description for the appropriate MMU */
/* type in Volume III of this book for the exact translation mechanism */

endfunction AddressTranslation
```

3.2.2.2.3 LoadMemory

The LoadMemory function loads a value from memory.

This action uses cache and main memory as specified in both the Cacheability and Coherency Attribute (*CCA*) and the access (*IorD*) to find the contents of *AccessLength* memory bytes, starting at physical location *pAddr*. The data is returned in a fixed-width naturally aligned memory element (*MemElem*). The low-order 2 (or 3) bits of the address and the *AccessLength* indicate which of the bytes within *MemElem* need to be passed to the processor. If the memory access type of the reference is *uncached*, only the referenced bytes are read from memory and marked as valid within the memory element. If the access type is *cached* but the data is not present in cache, an implementation-specific *size* and *alignment* block of memory is read and loaded into the cache to satisfy a load reference. At a minimum, this block is the entire memory element.

Figure 3.18 LoadMemory Pseudocode Function

```
MemElem ← LoadMemory (CCA, AccessLength, pAddr, vAddr, IorD)

/* MemElem: Data is returned in a fixed width with a natural alignment. The */
/*           width is the same size as the CPU general-purpose register, */
/*           32 or 64 bits, aligned on a 32- or 64-bit boundary, */
/*           respectively. */
/* CCA:      Cacheability&CoherencyAttribute=method used to access caches */
/*           and memory and resolve the reference */

/* AccessLength: Length, in bytes, of access */
/* pAddr:      physical address */
/* vAddr:      virtual address */
/* IorD:      Indicates whether access is for Instructions or Data */

endfunction LoadMemory
```

3.2.2.2.4 StoreMemory

The StoreMemory function stores a value to memory.

The specified data is stored into the physical location *pAddr* using the memory hierarchy (data caches and main memory) as specified by the Cacheability and Coherency Attribute (*CCA*). The *MemElem* contains the data for an aligned, fixed-width memory element (a word for 32-bit processors, a doubleword for 64-bit processors), though only the bytes that are actually stored to memory need be valid. The low-order two (or three) bits of *pAddr* and the *AccessLength* field indicate which of the bytes within the *MemElem* data should be stored; only these bytes in memory will actually be changed.

Figure 3.19 StoreMemory Pseudocode Function

```
StoreMemory (CCA, AccessLength, MemElem, pAddr, vAddr)

/* CCA:      Cacheability&Coherency Attribute, the method used to access */
/*           caches and memory and resolve the reference. */
/* AccessLength: Length, in bytes, of access */
/* MemElem:   Data in the width and alignment of a memory element. */
/*           The width is the same size as the CPU general */
/*           purpose register, either 4 or 8 bytes, */
/*           aligned on a 4- or 8-byte boundary. For a */
/*           partial-memory-element store, only the bytes that will be */
/*           stored must be valid. */
/* pAddr:     physical address */
/* vAddr:     virtual address */

endfunction StoreMemory
```

3.2.2.2.5 Prefetch

The Prefetch function prefetches data from memory.

Prefetch is an advisory instruction for which an implementation-specific action is taken. The action taken may increase performance but must not change the meaning of the program or alter architecturally visible state.

Figure 3.20 Prefetch Pseudocode Function

```
Prefetch (CCA, pAddr, vAddr, DATA, hint)

/* CCA:      Cacheability&Coherency Attribute, the method used to access */
/*           caches and memory and resolve the reference. */
/* pAddr:     physical address */
/* vAddr:     virtual address */
/* DATA:     Indicates that access is for DATA */
/* hint:      hint that indicates the possible use of the data */

endfunction Prefetch
```

Table 3.1 lists the data access lengths and their labels for loads and stores.

Table 3.1 AccessLength Specifications for Loads/Stores

AccessLength Name	Value	Meaning
DOUBLEWORD	7	8 bytes (64 bits)

Table 3.1 AccessLength Specifications for Loads/Stores

AccessLength Name	Value	Meaning
SEPTIBYTE	6	7 bytes (56 bits)
SEXTIBYTE	5	6 bytes (48 bits)
QUINTIBYTE	4	5 bytes (40 bits)
WORD	3	4 bytes (32 bits)
TRIPLEBYTE	2	3 bytes (24 bits)
HALFWORD	1	2 bytes (16 bits)
BYTE	0	1 byte (8 bits)

3.2.2.2.6 SyncOperation

The SyncOperation function orders loads and stores to synchronize shared memory.

This action makes the effects of the synchronizable loads and stores indicated by *stype* occur in the same order for all processors.

Figure 3.21 SyncOperation Pseudocode Function

```

SyncOperation(stype)

    /* stype: Type of load/store ordering to perform. */

    /* Perform implementation-dependent operation to complete the */
    /* required synchronization operation */

endfunction SyncOperation

```

3.2.2.3 Floating Point Functions

The pseudocode shown in below specifies how the unformatted contents loaded or moved to CP1 registers are interpreted to form a formatted value. If an FPR contains a value in some format, rather than unformatted contents from a load (uninterpreted), it is valid to interpret the value in that format (but not to interpret it in a different format).

3.2.2.3.1 ValueFPR

The ValueFPR function returns a formatted value from the floating point registers.

Figure 3.22 ValueFPR Pseudocode Function

```

value ← ValueFPR(fpr, fmt)

    /* value: The formatted value from the FPR */

    /* fpr:   The FPR number */
    /* fmt:   The format of the data, one of: */
    /*       S, D, W, L, PS, */
    /*       OB, QH, */
    /*       UNINTERPRETED_WORD, */
    /*       UNINTERPRETED_DOUBLEWORD */
    /* The UNINTERPRETED values are used to indicate that the datatype */
    /* is not known as, for example, in SWC1 and SDC1 */

```

```

case fmt of
  S, W, UNINTERPRETED_WORD:
    valueFPR ← UNPREDICTABLE32 || FPR[fpr]31..0

  D, UNINTERPRETED_DOUBLEWORD:
    if (FP32RegistersMode = 0)
      if (fpr0 ≠ 0) then
        valueFPR ← UNPREDICTABLE
      else
        valueFPR ← FPR[fpr+1]31..0 || FPR[fpr]31..0
      endif
    else
      valueFPR ← FPR[fpr]
    endif

  L, OB, QH:
    if (FP32RegistersMode = 0) then
      valueFPR ← UNPREDICTABLE
    else
      valueFPR ← FPR[fpr]
    endif

  DEFAULT:
    valueFPR ← UNPREDICTABLE

endcase
endfunction ValueFPR

```

The pseudocode shown below specifies the way a binary encoding representing a formatted value is stored into CP1 registers by a computational or move operation. This binary representation is visible to store or move-from instructions. Once an FPR receives a value from the StoreFPR(), it is not valid to interpret the value with ValueFPR() in a different format.

3.2.2.3.2 StoreFPR

Figure 3.23 StoreFPR Pseudocode Function

```

StoreFPR (fpr, fmt, value)

/* fpr:   The FPR number */
/* fmt:   The format of the data, one of: */
/*        S, D, W, L, PS, */
/*        OB, QH, */
/*        UNINTERPRETED_WORD, */
/*        UNINTERPRETED_DOUBLEWORD */
/* value: The formatted value to be stored into the FPR */

/* The UNINTERPRETED values are used to indicate that the datatype */
/* is not known as, for example, in LWC1 and LDC1 */

case fmt of
  S, W, UNINTERPRETED_WORD:
    FPR[fpr] ← UNPREDICTABLE32 || value31..0

  D, UNINTERPRETED_DOUBLEWORD:

```

```

        if (FP32RegistersMode = 0)
            if (fpr0 ≠ 0) then
                UNPREDICTABLE
            else
                FPR[fpr] ← UNPREDICTABLE32 || value31..0
                FPR[fpr+1] ← UNPREDICTABLE32 || value63..32
            endif
        else
            FPR[fpr] ← value
        endif

L, OB, QH:
    if (FP32RegistersMode = 0) then
        UNPREDICTABLE
    else
        FPR[fpr] ← value
    endif

endcase

endfunction StoreFPR

```

3.2.2.3.3 CheckFPEException

The pseudocode shown below checks for an enabled floating point exception and conditionally signals the exception.

Figure 3.24 CheckFPEException Pseudocode Function

```

CheckFPEException()

/* A floating point exception is signaled if the E bit of the Cause field is a 1 */
/* (Unimplemented Operations have no enable) or if any bit in the Cause field */
/* and the corresponding bit in the Enable field are both 1 */

if ( (FCSR17 = 1) or
      ((FCSR16..12 and FCSR11..7) ≠ 0) ) then
    SignalException(FloatingPointException)
endif

endfunction CheckFPEException

```

3.2.2.3.4 FPConditionCode

The FPConditionCode function returns the value of a specific floating point condition code.

Figure 3.25 FPConditionCode Pseudocode Function

```

tf ← FPConditionCode(cc)

/* tf: The value of the specified condition code */

/* cc: The Condition code number in the range 0..7 */

if cc = 0 then
    FPConditionCode ← FCSR23
else
    FPConditionCode ← FCSR24+cc

```

```

endif

endfunction FPConditionCode

```

3.2.2.3.5 SetFPConditionCode

The SetFPConditionCode function writes a new value to a specific floating point condition code.

Figure 3.26 SetFPConditionCode Pseudocode Function

```

SetFPConditionCode(cc, tf)
  if cc = 0 then
    FCSR ← FCSR31..24 || tf || FCSR22..0
  else
    FCSR ← FCSR31..25+cc || tf || FCSR23+cc..0
  endif

endfunction SetFPConditionCode

```

3.2.2.4 Pseudocode Functions Related to Sign and Zero Extension

3.2.2.4.1 Sign extension and zero extension in pseudocode

Much pseudocode uses a generic function `sign_extend` without specifying from what bit position the extension is done, when the intention is obvious. E.g. `sign_extend(immediate16)` or `sign_extend(dispatch9)`.

However, sometimes it is necessary to specify the bit position. For example, `sign_extend(temp31..0)` or the more complicated `(offset15)GPRLEN-(16+2) || offset || 02`.

The explicit notation `sign_extend.nbits(val)` or `sign_extend(val, nbits)` is suggested as a simplification. They say to sign extend as if an `nbits`-sized signed integer. The width to be sign extended to is usually apparent by context, and is usually `GPRLEN`, 32 or 64 bits. The previous examples then become.

```

sign_extend(temp31..0)
= sign_extend.32(temp)

```

and

```

(offset15)GPRLEN-(16+2) || offset || 02
= sign_extend.16(offset) << 2

```

Note that `sign_extend.N(value)` extends from bit position `N-1`, if the bits are numbered `0..N-1` as is typical.

The explicit notations `sign_extend.nbits(val)` or `sign_extend(val, nbits)` is used as a simplification. These notations say to sign extend as if an `nbits`-sized signed integer. The width to be sign extended to is usually apparent by context, and is usually `GPRLEN`, 32 or 64 bits.

Figure 3.27 sign_extend Pseudocode Functions

```

sign_extend.nbits(val) = sign_extend(val, nbits) /* syntactic equivalents */

function sign_extend(val, nbits)
  return (valnbits-1)GPRLEN-nbits || valnbits-1..0
end function

```

The earlier examples can be expressed as

```
(offset15)GPRLEN-(16+2) || offset || 02
= sign_extend.16(offset) << 2)
```

and

```
sign_extend(temp31..0)
= sign_extend.32(temp)
```

Similarly for `zero_extension`, although zero extension is less common than sign extension in the MIPS ISA.

Floating point may use notations such as `zero_extend.fmt` corresponding to the format of the FPU instruction. E.g. `zero_extend.S` and `zero_extend.D` are equivalent to `zero_extend.32` and `zero_extend.64`.

Existing pseudocode may use any of these, or other, notations.

3.2.2.4.2 memory_address

The pseudocode function `memory_address` performs mode-dependent address space wrapping for compatibility between MIPS32 and MIPS64. It is applied to all memory references. It may be specified explicitly in some places, particularly for new memory reference instructions, but it is also declared to apply implicitly to all memory references as defined below. In addition, certain instructions that are used to calculate effective memory addresses but which are not themselves memory accesses specify `memory_address` explicitly in their pseudocode.

Figure 3.28 memory_address Pseudocode Function

```
function memory_address(ea)
  if User mode and Status.UX = 0 then return sign_extend.32(ea)
  /* Preliminary proposal to wrap privileged mode addresses */
  if Supervisor mode and Status.SX = 0 then return sign_extend.32(ea)
  if Kernel mode and Status.KX = 0 then return sign_extend.32(ea)
  /* if Hardware Page Table Walking, then wrap in same way as Kernel/VZ Root */
  return ea
end function
```

On a 32-bit CPU, `memory_address` returns its 32-bit effective address argument unaffected.

On a 64-bit processor, `memory_address` optionally truncates a 32-bit address by sign extension. It discards carries that may have propagated from the lower 32-bits to the upper 32-bits that would cause minor differences between MIPS32 and MIPS64 execution. It is used in certain modes¹ on a MIPS64 CPU where strict compatibility with MIPS32 is required. This behavior was and continues to be described in a section of Volume III of the MIPS ARM². However, the behavior was not formally described in pseudocode functions prior to Release 6.

In addition to the use of `memory_address` for all memory references (including load and store instructions, LL/SC), Release 6 extends this behavior to control transfers (branch and call instructions), and to the PC-relative address calculation instructions (ADDIUPC, AUIPC, ALUIPC). In newer instructions the function is explicit in the pseudocode.

Implicit address space wrapping for all instruction fetches is described by the following pseudocode fragment which should be considered part of instruction fetch:

-
1. Currently, if in User/Supervisor/Kernel mode and Status.UX/SX/KX=0.
 2. E.g. see section named “Special Behavior for Data References in User Mode with Status_{UX}=0”, in the MIPS(r) Architecture Reference Manual Volume III, the MIPS64(R) and microMIPS64(tm) Privileged Resource Architecture, e.g. in section 4.11 of revision 5.03, or section 4.9 of revision 1.00.

Figure 3.29 Instruction Fetch Implicit memory_address Wrapping

```

PC ← memory_address( PC )
( instruction_data, length ) ← instruction_fetch( PC )
/* decode and execute instruction */

```

Implicit address space wrapping for all data memory accesses is described by the following pseudocode, which is inserted at the top of the AddressTranslation pseudocode function:

Figure 3.30 AddressTranslation implicit memory_address Wrapping

```

(pAddr, CCA) ← AddressTranslation( vAddr, IorD, LorS )
vAddr ← memory_address( vAddr )

```

In addition to its use in instruction pseudocode,

3.2.2.5 Miscellaneous Functions

This section lists miscellaneous functions not covered in previous sections.

3.2.2.5.1 SignalException

The SignalException function signals an exception condition.

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 3.31 SignalException Pseudocode Function

```

SignalException(Exception, argument)

/* Exception:    The exception condition that exists. */
/* argument:    A exception-dependent argument, if any */

endfunction SignalException

```

3.2.2.5.2 SignalDebugBreakpointException

The SignalDebugBreakpointException function signals a condition that causes entry into Debug Mode from non-Debug Mode.

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 3.32 SignalDebugBreakpointException Pseudocode Function

```

SignalDebugBreakpointException()

endfunction SignalDebugBreakpointException

```

3.2.2.5.3 SignalDebugModeBreakpointException

The SignalDebugModeBreakpointException function signals a condition that causes entry into Debug Mode from Debug Mode (i.e., an exception generated while already running in Debug Mode).

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 3.33 SignalDebugModeBreakpointException Pseudocode Function

```
SignalDebugModeBreakpointException()

endfunction SignalDebugModeBreakpointException
```

3.2.2.5.4 NullifyCurrentInstruction

The NullifyCurrentInstruction function nullifies the current instruction.

The instruction is aborted, inhibiting not only the functional effect of the instruction, but also inhibiting all exceptions detected during fetch, decode, or execution of the instruction in question. For branch-likely instructions, nullification kills the instruction in the delay slot of the branch likely instruction.

Figure 3.34 NullifyCurrentInstruction PseudoCode Function

```
NullifyCurrentInstruction()

endfunction NullifyCurrentInstruction
```

3.2.2.5.5 NotWordValue

The NotWordValue function returns a boolean value that determines whether the 64-bit value contains a valid word (32-bit) value. Such a value has bits 63..32 equal to bit 31.

Figure 3.35 NotWordValue Pseudocode Function

```
result ← NotWordValue(value)

/* result:    True if the value is not a correct sign-extended word value; */
/*           False otherwise */

/* value:     A 64-bit register value to be checked */

NotWordValue ← value63..32 ≠ (value31)32

endfunction NotWordValue
```

3.2.2.5.6 PolyMult

The PolyMult function multiplies two binary polynomial coefficients.

Figure 3.36 PolyMult Pseudocode Function

```
PolyMult(x, y)
    temp ← 0
    for i in 0 .. 31
        if xi = 1 then
            temp ← temp xor (y(31-i)..0 || 0i)
        endif
    endfor

    PolyMult ← temp

endfunction PolyMult
```

3.3 Op and Function Subfield Notation

In some instructions, the instruction subfields *op* and *function* can have constant 5- or 6-bit values. When reference is made to these instructions, uppercase mnemonics are used. For instance, in the floating point ADD instruction, *op*=COP1 and *function*=ADD. In other cases, a single field has both fixed and variable subfields, so the name contains both upper- and lowercase characters.

3.4 FPU Instructions

In the detailed description of each FPU instruction, all variable subfields in an instruction format (such as *fs*, *ft*, *immediate*, and so on) are shown in lowercase. The instruction name (such as ADD, SUB, and so on) is shown in uppercase.

For the sake of clarity, an alias is sometimes used for a variable subfield in the formats of specific instructions. For example, *rs=base* in the format for load and store instructions. Such an alias is always lowercase since it refers to a variable subfield.

Bit encodings for mnemonics are given in Volume I, in the chapters describing the CPU, FPU, MDMX, and MIPS16e instructions.

See “Op and Function Subfield Notation” on page 51 for a description of the *op* and *function* subfields.

Instruction Formats

This chapter defines the formats of microMIPS instructions. The microMIPS variable-length encoding comprises 16-bit and 32-bit wide instructions. The 6-bit major opcode is left-aligned within the instruction encoding. Instructions can have 0 to 4 register fields. For 32-bit instructions, the register field width is 5 bits, while for most 16-bit instructions, the register field width is 3 bits, utilizing instruction-specific register encoding. All 5-bit register fields are located at a constant position within the instruction encoding.

The immediate field is right-aligned in the following instructions:

- some 16-bit instructions with 3-bit register fields
- 32-bit instructions with 16-bit or 26-bit immediate field

The name ‘immediate field’ as used here includes the address offset field for branches and load/store instructions as well as the jump target field.

Other instruction-specific fields are typically located between the immediate and minor opcode fields. Instructions that have multiple “other” fields are listed in alphabetical order according to the name of the field, with the first name of the order located at the lower bit position. An empty bit field that is not explicitly shown in the instruction format is located next to the minor opcode field.

[Figure 4.1](#) and [Figure 4.2](#) show the 16-bit and 32-bit instruction formats.

Figure 4.1 16-Bit Instruction Formats

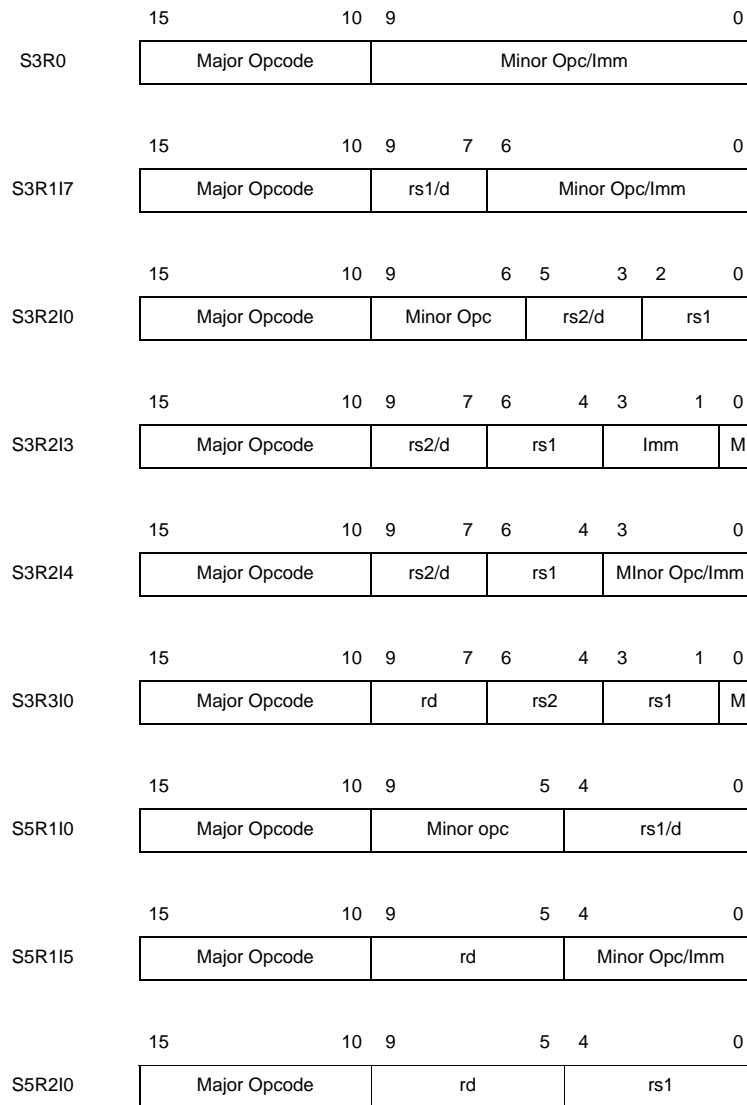


Figure 4.2 32-Bit Instruction Formats

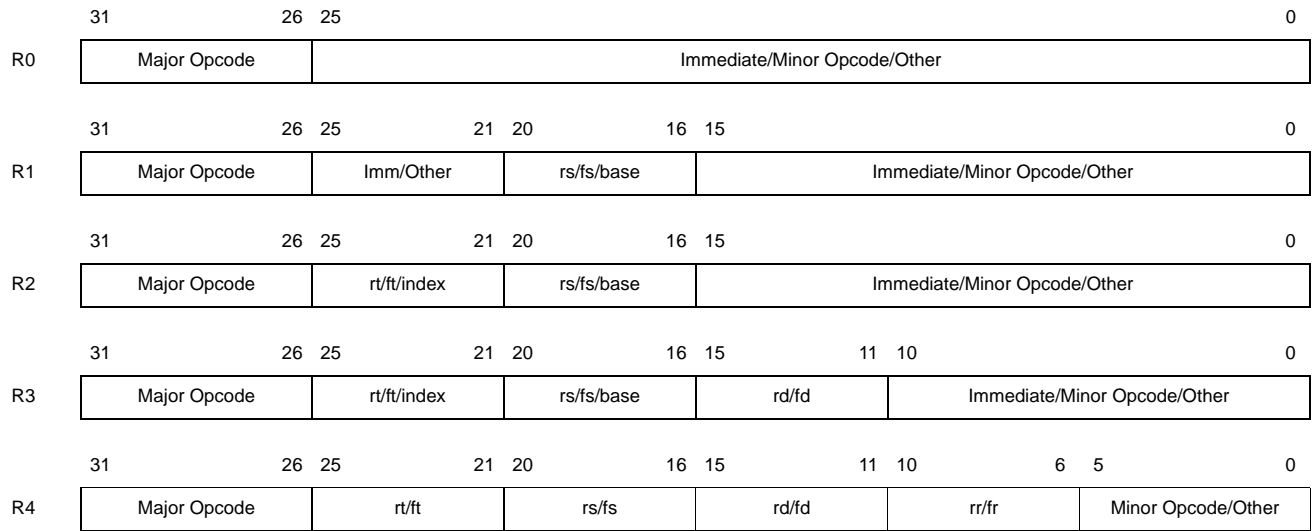




Figure 4.3 Immediate Fields within 32-Bit Instructions


32-bit instruction formats with 26-bit immediate fields:



32-bit instruction formats with 16-bit immediate fields:



32-bit instruction formats with 12-bit immediate fields:



Instruction Formats

The instruction size can be completely derived from the major opcode. For 32-bit instructions, the major opcode also defines the position of the minor opcode field and whether or not the immediate field is right-aligned.

Instructions formats are named according to the number of the register fields and the size of the immediate field. The names have the structure R<x>I<y>. For example, an instruction based on the format R2I16 has 2 register fields and a 16-bit immediate field.

[Table 4.1](#) shows all formats. The 16-bit formats refer to either 3-bit or 5-bit register fields. To visualize this, a 16-bit format name starts with the prefix S3 or S5 respectively.

Table 4.1 microMIPS Opcode Formats

32-bit Instruction Formats (existing instructions)	32-bit Instruction Formats (additional format(s) for new instructions)	16-bit Instruction Formats
R0I0	R2I12	S3R0I0
R0I8		S3R0I10
R0I16		S3R1I7
R0I26		S3R2I0
R1I0		S3R2I3
R1I2		S3R2I4
R1I7		S3R3I1
R1I8		S5R1I0
R1I10		S5R1I4
R1I16		S5R2I0
R2I0		
R2I2		
R2I3		
R2I4		
R2I5		
R2I10		
R2I16		
R3I0		
R3I3		
R4I0		

4.1 Instruction Stream Organization and Endianness

16-bit instructions are placed within the 32-bit (or 64-bit) memory element according to system endianness.

- On a 32-bit processor in big-endian mode, the first instruction is read from bits 31..16, and the second instruction is read from bits 15..0.
- On a 32-bit processor in little-endian mode, the first instruction is read from bits 15..0, and the second instruction is read from bits 31..16.

The above rule also applies to the halfwords of 32-bit instructions. This means that a 32-bit instruction is not treated as a word data type; instead, the halfwords are treated in the same way as individual 16-bit instructions. The halfword containing the major opcode is always the first in the sequence.

Example:

SRL r1, r1, 7	binary opcode fields:	000000 00001 00001 00111 00001 000000
	hex representation:	0021 3840
	Address:	3 2 1 0
Little Endian:	Data:	38 40 00 21
	Address:	0 1 2 3
Big Endian:	Data:	00 21 38 40

Instructions are placed in memory such that they are in-order with respect to the address.

microMIPS Instruction Set

This chapter lists all microMIPS encoded instructions, sorted into 16-bit and 32-bit categories.

In the 16-bit category:

- Frequent MIPS64 instructions and macros, re-encoded as 16-bit. Register and immediate fields are reduced in size by using encodings of frequently occurring values.

In the 32-bit category:

- All MIPS64 instructions, including all application-specific extensions except MIPS16e, re-encoded: MIPS64, MDMX ASE, MIPS-3D ASE, MIPS DSP ASE, MIPS MT ASE, and SmartMIPS ASE.
- Opcode space for user-defined instructions (UDIs).
- New instructions designed primarily to reduce code size.

To differentiate between 16-bit and 32-bit encoded instructions, the instruction mnemonic can be optionally extended with the suffix “16” or “32” respectively. This suffix is placed at the end of the instruction before the first ‘.’ if there is one. For example:

ADD16, ADD32, ADD32.PS, DADDIU32

If these suffixes are omitted, the assembler automatically chooses the smallest instruction size.

For each instruction, the tables in this chapter provide all necessary information about the bit fields. The formats of the instructions are defined in [Chapter 4, “Instruction Formats” on page 53](#). Together with the major and minor opcode encodings, which can be derived from the tables in [Chapter 7, “Opcode Map” on page 409](#), the complete instruction encoding is provided.

Most register fields have a width of 5 bits. 5-bit register fields use linear encoding ($r0=’00000’$, $r1=’00001’$, etc.). For 16-bit instructions, whose register field size is variable, the register field width is explicitly stated in the instruction table ([Table 5.1](#) and [Table 5.2](#)), and the individual register and immediate encodings are shown in [Table 5.3](#). The ‘other fields’ are defined by the respective column, with the order of these fields in the instruction encoding defined by the order in the tables.

5.1 16-Bit Category

5.1.1 Frequent MIPS64 Instructions

These are frequent MIPS64 instructions with reduced register and immediate fields containing frequently used registers and immediate values.

MOVE is a very frequent instruction. It therefore supports full 5-bit unrestricted register fields for maximum efficiency. In fact, MOVE used to be a simplified macro of an existing MIPS64 instruction.

There are 2 variants of the LW and SW instructions. One variant implicitly uses the SP register to allow for a larger offset field. The value in the offset field is shifted left by 2 before it is added to the base address.

There are four variants of the ADDIU instruction:

1. A variant with one 5-bit register specifier that allows any GPR to be the source and destination register
2. A variant that uses the stack pointer as the implicit source and destination register
3. A variant that has separate 3-bit source and destination register specifiers
4. A variant that has the stack pointer as the implicit source register and one 3-bit destination register specifier

A 16-bit NOP instruction is needed because of the new 16-bit instruction alignment and the need in specific cases to align instructions on a 32-bit boundary. It can save code size as well. NOP is not shown in the table because it is realized as a macro (as is NEGU).

```
NOP16 = MOVE16 r0, r0
```

```
NEGU16 rt, rs = SUBU16 rt, r0, rs
```

Because microMIPS instructions are 16-bit aligned, the 16-bit branch instructions support 16-bit aligned branch target addresses. The offset field is left shifted by 1 before it is added to the PC.

The breakpoint instructions, BREAK and SDBBP, include a 16-bit variant that allows a breakpoint to be inserted at any instruction address without overwriting more than a single instruction.

The instructions in the following tables are pre-Release 6 instructions. Refer to [Section 2.7 “Release 6 of the MIPS Architecture”](#) to understand which instructions have been removed in Release 6.

Table 5.1 16-Bit Re-encoding of Frequent MIPS64 Instructions

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
ADDIUS5	POOL16D	5bit:1	4	5		0	1	Add Immediate Unsigned Word Same Register
ADDIUSP	POOL16D	0	9	0		0	1	Add Immediate Unsigned Word to Stack Pointer
ADDIUR2	POOL16E	2	3	3		0	1	Add Immediate Unsigned Word Two Registers
ADDIUR1SP	POOL16E	1	6	3		0	1	Add Immediate Unsigned Word One Registers and Stack Pointer
ADDU16	POOL16A	3	0	3		0	1	Add Unsigned Word

Table 5.1 16-Bit Re-encoding of Frequent MIPS64 Instructions (Continued)

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
AND16	POOL16C	2	0	3		0	4	AND
ANDI16	ANDI16	2	4	3		0	0	AND Immediate
B16	B16	0	10			0	0	Branch
BREAK16	POOL16C	0	0		4	0	6	Cause Breakpoint Exception
JALR16	POOL16C	1	0	5		0	5	Jump and Link Register, 32-bit delay-slot
JALRS16	POOL16C	1	0	5		0	5	Jump and Link Register, 16-bit delay-slot
JR16	POOL16C	1	0	5		0	5	Jump Register
LBU16	LBU16	2	4	3		0	0	Load Byte Unsigned
LHU16	LHU16	2	4	3		0	0	Load Halfword
LI16	LI16	1	7	3		0	0	Load Immediate
LW16	LW16	2	4	3		0	0	Load Word
LWGP	LWGP16	1	7	3		0	0	Load Word GP
LWSP	LWSP16	5bit:1	5	5		0	0	Load Word SP
MFHI16	POOL16C	1	0	5		0	5	Move from HI Register
MFLO16	POOL16C	1	0	5		0	5	Move from LO Register
MOVE16	MOVE16	2	0	5		0	0	Move
NOT16	POOL16C	2	0	3		0	4	NOT
OR16	POOL16C	2	0	3		0	4	OR
SB16	SB16	2	4	3		0	0	Store Byte
SDBBP16	POOL16C	0	0		4	0	6	Cause Debug Breakpoint Exception
SH16	SH16	2	4	3		0	0	Store Halfword
SLL16	POOL16B	2	3	3		0	1	Shift Word Left Logical
SRL16	POOL16B	2	3	3		0	1	Shift Word Right Logical
SUBU16	POOL16A	3	0	3		0	1	Sub Unsigned
SW16	SW16	2	4	3		0	0	Store Word
SWSP	SWSP16	5bit:1	5	5		0	0	Store Word SP
XOR16	POOL16C	2	0	3		0	4	XOR

5.1.2 Frequent MIPS64 Instruction Sequences

These 16-bit instructions are equivalent to frequently-used short sequences of MIPS64 instructions. The instruction-specific register and immediate value selection are shown in [Table 5.3](#).

Table 5.2 16-Bit Re-encoding of Frequent MIPS64 Instruction Sequences

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
BEQZ16	BEQZ16	1	7	3		0	0	Branch on Equal Zero
BNEZ16	BNEZ16	1	7	3		0	0	Branch on Not Equal Zero
JRADDIUSP	POOL16C	0	5				5	Jump Register; ADDIU SP
JRC	POOL16C	1	0	5		0	5	Jump Register Compact
LWM16	POOL16C	0	4		2	0	4	Load Word Multiple
MOVEP	POOL16C	3 (encoded)	0	3(encoded)		0	1	Move Register Pair
SWM16	POOL16C	0	4		2	0	4	Store Word Multiple

5.1.3 Instruction-Specific Register Specifiers and Immediate Field Encodings

Table 5.3 Instruction-Specific Register Specifiers and Immediate Field Values

Instruction	Number of Register Fields	Immediate Field Size (bit)	Register 1 Decoded Value	Register 2 Decoded Value	Register 3 Decoded Value	Immediate Field Decoded Value
ADDIUS5	5bit:1	4	rd: 5 bit field			-8..0..7
ADDIUSP	0	9				$(-258..-3, 2..257) \ll 2$
ADDIUR2	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		-1, 1, 4, 8, 12, 16, 20, 24
ADDIUR1SP	1	6	rd:2-7,16, 17			$(0..63) \ll 2$
ADDU16	3	0	rs1:2-7,16, 17	rs2:2-7,16, 17	rd:2-7,16, 17	
AND16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
ANDI16	2	4	rs1:2-7,16, 17	rd:2-7,16, 17		1, 2, 3, 4, 7, 8, 15, 16, 31, 32, 63, 64, 128, 255, 32768, 65535
B16	0	10				$(-512..511) \ll 1$
BEQZ16	1	7	rs1:2-7,16, 17			$(-64..63) \ll 1$
BNEZ16	1	7	rs1:2-7,16, 17			$(-64..63) \ll 1$
BREAK16	0	4				0..15
JALR16	5bit:1	0	rs1:5 bit field			
JALRS16	5bit:1	0	rs1:5 bit field			
JRADDIUSP	0	5				$(0..31) \ll 2$
JR16	5bit:1	0	rs1:5 bit field			
JRC	5bit:1	0	rs1:5 bit field			
LBU16	2	4	rb:2-7,16,17	rd:2-7,16, 17		-1,0..14
LHU16	2	4	rb:2-7,16,17	rd:2-7,16, 17		$(0..15) \ll 1$
LI16	1	7	rd:2-7,16, 17			-1,0..126
LW16	2	4	rb:2-7,16,17	rd:2-7,16, 17		$(0..15) \ll 2$
LWM16	2bit list:1	4				$(0..15) \ll 2$
LWGP	1	7	rd:2-7,16,17			$(-64..63) \ll 2$
LWSP	5bit:1	5	rd:5-bit field			$(0..31) \ll 2$
MFHI16	5bit:1	0	rd:5-bit field			
MFLO16	5bit:1	0	rd:5-bit field			
MOVE16	5bit:2	0	rd:5-bit field	rs1:5-bit field		
MOVEP	3	0	rd, re: (5,6),(5,7),(6,7), (4,21),(4,22),(4, 5),(4,6),(4,7)	rt:0,2,7,16-20	rs:0,2,7,16-20	
NOT16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
OR16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
SB16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		0..15

Table 5.3 Instruction-Specific Register Specifiers and Immediate Field Values (Continued)

Instruction	Number of Register Fields	Immediate Field Size (bit)	Register 1 Decoded Value	Register 2 Decoded Value	Register 3 Decoded Value	Immediate Field Decoded Value
SDBBP16	0	0				0..15
SH16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		(0..15) << 1
SLL16	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		1..8 (see encoding tables)
SRL16	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		1..8 (see encoding tables)
SUBU16	3	0	rs1:2-7,16, 17	rs2:2-7,16, 17	rd:2-7,16, 17	
SW16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		(0..15) << 2
SWSP	5bit:1	5	rs1: 5 bit field			(0..31) << 2
SWM16	2 bit list:1	4				(0..15)<<2
XOR16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		

5.2 16-bit Instruction Register Set

Many of the 16-bit instructions use 3-bit register specifiers in their binary encodings. The register set used for most of these 3-bit register specifiers is listed in [Table 5.5](#). The register set used for SB16, SH16, SW16 source register is listed in [Table 5.5](#). These register sets are a true subset of the register set available in 32-bit mode; the 3-bit register specifiers can directly access 8 of the 32 registers available in 32-bit mode (which uses 5-bit register specifiers).

In addition, specific instructions in the 16-bit instruction set implicitly reference the stack pointer register (*sp*), global pointer register (*gp*), the return address register (*ra*), the integer multiplier/divider output registers (*HI/LO*) and the program counter (*PC*). Of these, [Table 5.6](#) lists *sp*, *gp* and *ra*. [Table 5.7](#) lists the microMIPS special-purpose registers, including *PC*, *HI* and *LO*.

The microMIPS also contains some 16-bit instructions that use 5-bit register specifiers. Such 16-bit instructions provide access to all 32 general-purpose registers.

Table 5.4 16-Bit Instruction General-Purpose Registers - \$2-\$7, \$16, \$17

16-Bit Register Encoding ¹	32-Bit MIPS Register Encoding ²	Symbolic Name (From <i>ArchDefs.h</i>)	Description
0	16	s0	General-purpose register
1	17	s1	General-purpose register
2	2	v0	General-purpose register
3	3	v1	General-purpose register
4	4	a0	General-purpose register
5	5	a1	General-purpose register
6	6	a2	General-purpose register

Table 5.4 16-Bit Instruction General-Purpose Registers - \$2-\$7, \$16, \$17 (Continued)

16-Bit Register Encoding ¹	32-Bit MIPS Register Encoding ²	Symbolic Name (From <i>ArchDefs.h</i>)	Description
7	7	a3	General-purpose register

1. “0-7” correspond to the register’s 16-bit binary encoding and show how that encoding relates to the MIPS registers. “0-7” never refer to the registers, except within the binary microMIPS instructions. From the assembler, only the MIPS names (\$16, \$17, \$2, etc.) or the symbolic names (s0, s1, v0, etc.) refer to the registers. For example, to access register number 17 in the register file, the programmer references \$17 or s1, even though the micro-MIPS binary encoding for this register is 001.
2. General registers not shown in the above table are not accessible through the 16-bit instruction using 3-bit register specifiers. The Move instruction can access all 32 general-purpose registers.

Table 5.5 SB16, SH16, SW16 Source Registers - \$0, \$2-\$7, \$17

16-Bit Register Encoding ¹	32-Bit MIPS Register Encoding ²	Symbolic Name (From <i>ArchDefs.h</i>)	Description
0	0	zero	Hard-wired Zero
1	17	s1	General-purpose register
2	2	v0	General-purpose register
3	3	v1	General-purpose register
4	4	a0	General-purpose register
5	5	a1	General-purpose register
6	6	a2	General-purpose register
7	7	a3	General-purpose register

1. “0-7” correspond to the register’s 16-bit binary encoding and show how that encoding relates to the MIPS registers. “0-7” never refer to the registers, except within the binary microMIPS instructions. From the assembler, only the MIPS names (\$16, \$17, \$2, etc.) or the symbolic names (s0, s1, v0, etc.) refer to the registers. For example, to access register number 17 in the register file, the programmer references \$17 or s1, even though the micro-MIPS binary encoding for this register is 001.
2. General registers not shown in the above table are not accessible through the 16-bit instructions using 3-bit register specifier. The Move instruction can access all 32 general-purpose registers.

Table 5.6 16-Bit Instruction Implicit General-Purpose Registers

16-Bit Register Encoding	32-Bit MIPS Register Encoding	Symbolic Name (From <i>ArchDefs.h</i>)	Description
Implicit	28	gp	Global pointer register
Implicit	29	sp	Stack pointer register
Implicit	31	ra	Return address register

Table 5.7 16-Bit Instruction Special-Purpose Registers

Symbolic Name	Purpose
PC	Program counter. The PC-relative ADDIU can access this register as an operand.
HI	Contains high-order word of multiply or divide result.
LO	Contains low-order word of multiply or divide result.

5.3 32-Bit Category

The instructions in the following tables are pre-Release 6 instructions. Refer to [Section 2.7 “Release 6 of the MIPS Architecture”](#) to understand which instructions have been removed in Release 6.

5.3.1 New 32-bit instructions

The following table lists the 32-bit instructions introduced in the microMIPS ISA. Only instructions introduced prior to Release 6 are included in this table. JALRS, JALRS.HB, JALS, and JALX have been removed in Release 6.

Table 5.8 32-bit Instructions introduced within microMIPS

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
ADDIUPC	ADDIUPC	1	23	3		0	0	ADDIU PC-Relative
BEQZC	POOL32I	2:5 bit	16	5			0	Branch on Equal to Zero, No Delay Slot
BNEZC	POOL32I	2:5 bit	16	5			0	Branch on Not Equal to Zero, No Delay Slot
JALRS	POOL32A	2:5 bit	0	5			16	Jump and Link Register, Short Delay Slot


Table 5.8 32-bit Instructions introduced within microMIPS (Continued)

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
JALRS.HB	POOL32A	2:5 bit	0	5			16	Jump and Link Register with Hazard Barrier, Short Delay Slot
JALS	JALS32	0	26				0	Jump and Link, Short Delay Slot
JALX	JALX		26	5		0	5	Jump and Link Exchange
LDM	POOL32B	1:5bit	12		5	0	4	Load DoubleWord Multiple
LDP	POOL32B	2:5 bit	12			0	4	Load DoubleWord Pair
LWP	POOL32B	2:5 bit	12		5	0	4	Load Word Pair
LWXS	POOL32A	3:5 bit	0	5	0	1	10	Load Word Indexed, Scale
LWM32	POOL32B	1:5bit	12		5	0	4	Load Word Multiple
SDM	POOL32B	1:5bit	12		5	0	4	Store DoubleWord Multiple
SDP	POOL32B	2:5 bit	12			0	4	Load DoubleWord Pair
SWP	POOL32B	2:5 bit	12			0	4	Load Word Pair
SWM32	POOL32B	1:5bits	12		5	0	4	Store Word Multiple

5.4 Instructions Specific to microMIPS

This section describes instructions unique to microMIPS.

Only instructions supported in Release 6 are provided. [Section 2.7, "Release 6 of the MIPS Architecture,"](#) lists instructions that have been added, removed and recoded in Release 6.



Format: ADDIUR1SP rd, decoded_immediate_value

microMIPS

Purpose: Add Immediate Unsigned Word One Register (16-bit instr size)

To add a constant to a 32-bit integer.

Description: $GPR[rd] \leftarrow GPR[29] + \text{zero_extend}(\text{immediate} \ll 2)$

The 6-bit *immediate* field is first shifted left by two bits and then zero-extended. This amount is added to the 32-bit value in GPR 29 and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← GPR[29] + zero_extend(immediate || 02)
GPR[rd] ← sign_extend(temp31..0)


```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDIUR2 rd, rs1, decoded_immediate_value

microMIPS

Purpose: Add Immediate Unsigned Word Two Registers (16-bit instr size)

To add a constant to a 32-bit integer.

Description: $GPR[rd] \leftarrow GPR[rs] + \text{sign_extend}(\text{decoded_immediate})$

The encoded immediate field is decoded to obtain the actual immediate value.

The decoded immediate value is sign-extended and then added to the 32-bit value in GPR *rs*, and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Table 5.9 Encoded and Decoded Values of the Immediate Field

Encoded Value of Instr _{3..1} (Decimal)	Encoded Value of Instr _{3..1} (Hex)	Decoded Value of Immediate (Decimal)	Decoded Value of Immediate (Hex)
0	0x0	1	0x0001
1	0x1	4	0x0004
2	0x2	8	0x0008
3	0x3	12	0x000c
4	0x4	16	0x0010
5	0x5	20	0x0014
6	0x6	24	0x0018
7	0x7	-1	0xffff

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← GPR[rs] + sign_extend(decoded_immediate)
GPR[rd] ← sign_extend(temp31..0)


```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDIUS5 rd, decoded_immediate_value

microMIPS

Purpose: Add Immediate Unsigned Word 5-Bit Register Select (16-bit instr size)

To add a constant to a 32-bit integer

Description: $GPR[rd] \leftarrow GPR[rd] + \text{sign_extend}(\text{immediate})$

The 4-bit *immediate* field is sign-extended and then added to the 32-bit value in GPR *rd*. The 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

The 5-bit register select allows this 16-bit instruction to use any of the 32 GPRs as the destination register.

No Integer Overflow exception occurs under any circumstances.

Table 5-1 Encoded and Decoded Values of Signed Immediate Field

Encoded Value of Instr _{4..1} (Decimal)	Encoded Value of Instr _{4..1} (Hex)	Decoded Value of Immediate (Decimal)	Decoded Value of Immediate (Hex)
0	0x0	0	0x0000
1	0x1	1	0x0001
2	0x2	2	0x0002
3	0x3	3	0x0003
4	0x4	4	0x0004
5	0x5	5	0x0005
6	0x6	6	0x0006
7	0x7	7	0x0007
8	0x8	-8	0xffff8
9	0x9	-7	0xffff9
10	0xa	-6	0xffffa
11	0xb	-5	0xffffb
12	0xc	-4	0xffffc
13	0xd	-3	0xffffd
14	0xe	-2	0xffffe
15	0xf	-1	0xfffff

Restrictions:

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```
if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← GPR[rd] + sign_extend(immediate)
GPR[rd] ← sign_extend(temp31..0)
```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDIUSP decoded_immediate_value

microMIPS

Purpose: Add Immediate Unsigned Word to Stack Pointer(16-bit instr size)

To add a constant to the stack pointer.

Description: $GPR[29] \leftarrow GPR[29] + \text{sign_extend}(\text{decoded_immediate} \ll 2)$

The encoded immediate field is decoded to obtain the actual immediate value.

The actual immediate value is first shifted left by two bits and then sign-extended. This amount is added to the 32-bit value in GPR 29, and the 32-bit arithmetic result is sign-extended and placed into GPR 29.

No Integer Overflow exception occurs under any circumstances.

Table 5.10 Encoded and Decoded Values of Immediate Field

Encoded Value of Instr _{9..1} (Decimal)	Encoded Value of Instr _{9..1} (Hex)	Decoded Value of Immediate (Decimal)	Decoded Value of Immediate (Hex)
0	0x0	256	0x0100
1	0x1	257	0x0101
2	0x2	2	0x0002
3	0x3	3	0x0003
...
254	0xfe	254	0x00fe
255	0xff	255	0x00ff
256	0x100	-256	0xff00
257	0x101	-255	0xff01
...
508	0x1fc	-4	0xfffc
509	0x1fd	-3	0xfffd
510	0x1fe	-258	0xfefe
511	0x1ff	-257	0xfeff

Restrictions:

If GPR 29 does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```
if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← GPR[29] + sign_extend(decoded immediate || 02)
GPR[29] ← sign_extend(temp31..0)
```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDU16 rd, rs, rt

microMIPS

Purpose: Add Unsigned Word (16-bit instr size)

To add 32-bit integers

Description: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs*, and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

If either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← GPR[rs] + GPR[rt]
GPR[rd] ← sign_extend(temp31..0)

```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.

POOL16C 010001	rt	rs	AND16 0001
6	3	3	4

Format: AND16 *rt*, *rs*

microMIPS

Purpose: And (16-bit instr size)

To do a bitwise logical AND

Description: $GPR[rt] \leftarrow GPR[rs] \text{ AND } GPR[rt]$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical AND operation. The result is placed into GPR *rt*.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Availability and Compatibility:


This instruction has been recoded for Release 6.

Operation:

$GPR[rt] \leftarrow GPR[rs] \text{ and } GPR[rt]$

Exceptions:

None



Format: ANDI16 rd, rs, decoded_immediate_value

microMIPS

Purpose: And Immediate (16-bit instr size)

To do a bitwise logical AND with a constant

Description: $GPR[rd] \leftarrow GPR[rs] \text{ AND } \text{zero_extend}(\text{decoded immediate})$

The encoded immediate field is decoded to obtain the actual immediate value

The decoded immediate is zero-extended to the left and combined with the contents of GPR rs in a bitwise logical AND operation. The result is placed into GPR rd.

Table 5-2 Encoded and Decoded Values of Immediate Field

Encoded Value of Instr _{3..0} (Decimal)	Encoded Value of Instr _{3..0} (Hex)	Decoded Value of Immediate (Decimal)	Decoded Value of Immediate (Hex)
0	0x0	128	0x80
1	0x1	1	0x1
2	0x2	2	0x2
3	0x3	3	0x3
4	0x4	4	0x4
5	0x5	7	0x7
6	0x6	8	0x8
7	0x7	15	0xf
8	0x8	16	0x10
9	0x9	31	0x1f
10	0xa	32	0x20
11	0xb	63	0x3f
12	0xc	64	0x40
13	0xd	255	0xff
14	0xe	32768	0x8000
15	0xf	65535	0xffff

Restrictions:


The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ and } \text{zero_extend}(\text{decoded immediate})$

Exceptions:

None



Format: BC16 offset

microMIPS Release 6

Purpose: Unconditional Branch Compact (16-bit instr size)

To do an unconditional branch

Description: branch

A 11-bit signed offset (the 10-bit *offset* field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself) to form a PC-relative effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed when the branch is taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Operation:


```
target_offset ← sign_extend(offset || 01)
PC ← PC + 2 + target_offset
```

Exceptions:

None

Programming Notes:

With the 11-bit signed instruction offset, the branch range is ± 1 Kbytes. Use jump (JRC16 or JIC) or 32-bit branch instructions to branch to addresses outside this range.



Format: BEQZC16 rs, offset

microMIPS Release 6

Purpose: Branch on Equal to Zero Compact (16-bit instr size)

To compare a GPR to zero then do a PC-relative conditional branch

Description: if GPR[rs] = 0 then branch

A 8-bit signed offset (the 7-bit *offset* field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself) to form a PC-relative effective target address.

If the contents of GPR *rs* equals zero, branch to the effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed if the branch is taken.

Restrictions:

The 3-bit register field can only specify GPRs \$2-\$7, \$16, \$17.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Operation:

```

target_offset ← sign_extend(offset || 0)
condition ← (GPR[rs] == 0)
if condition then
    PC ← PC + target_offset
endif


```

Exceptions:

None

Programming Notes:

With the 8-bit signed instruction offset, the conditional branch range is ± 64 Bytes. Use 32-bit branch, jump (JRC16 or JIC) instructions to branch to addresses outside this range.



Format: BNEZC16 rs, offset

microMIPS Release 6

Purpose: Branch on Not Equal to Zero Compact (16-bit instr size)

To compare a GPR to zero then do a PC-relative conditional branch

Description: if GPR[rs] != 0 then branch

A 8-bit signed offset (the 7-bit *offset* field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

If the contents of GPR *rs* does not equal zero, branch to the effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed if the branch is taken.

Restrictions:

The 3-bit register field can only specify GPRs \$2-\$7, \$16, \$17.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Operation:

```

target_offset ← sign_extend(offset || 0)
condition ← (GPR[rs] != 0)
if condition then
    PC ← PC + target_offset
endif


```

Exceptions:

None

Programming Notes:

With the 8-bit signed instruction offset, the conditional branch range is ± 64 Bytes. Use 32-bit branch, jump (JRC16 or JIC) instructions to branch to addresses outside this range.



Format: BREAK16

microMIPS

Purpose: Breakpoint

To cause a Breakpoint exception

Description:

A breakpoint exception occurs, immediately and unconditionally transferring control to the exception handler. The *code* field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Restrictions:

None

Availability and Compatibility:


This instruction has been recoded for Release 6.

Operation:

SignalException(Breakpoint)

Exceptions:

Breakpoint



Format: JALRC16 rs

microMIPS Release 6

Purpose: Jump and Link Register Compact (16-bit instr size)

To execute a procedure call to an instruction address in a register

Description: $GPR[31] \leftarrow return_addr$, $PC \leftarrow GPR[rs]$

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

Place the return address link in GPR *r31*. The return link is the address of the first instruction following the branch, where execution continues after a procedure call.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors which implement MIPS64 and if the ISAMode bit of the target is MIPS64 (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS64 ISA, if the intended target ISAMode is MIPS64 (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump; that is, delay slot restrictions do not apply in Release 6.

Operation:


```

temp ← GPR[rs]
GPR[31] ← PC + 2
if Config3ISA = 1 then
  PC ← temp
else
  PC ← tempGPRLEN-1..1 || 0
  ISAMode ← temp0
endif

```

Exceptions:

None



Format: JRCADDIUSP decoded_immediate

microMIPS Release 6

Purpose: Jump Register Compact, Adjust Stack Pointer (16-bit)

To execute a branch to an instruction address in a register and adjust stack pointer

Description: $PC \leftarrow GPR[ra]; SP \leftarrow SP + \text{zero_extend}(\text{Immediate} \ll 2)$

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

The 5-bit *immediate* field is first shifted left by two bits and then zero-extended. This amount is then added to the 32-bit value of GPR 29 and the 32-bit arithmetic result is sign-extended and placed into GPR 29. No Integer Overflow exception occurs under any circumstances for the update of GPR 29.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors which implement MIPS64 and the ISAMode bit of the target address is MIPS64 (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS64 ISA, if the intended target ISAMode is MIPS64 (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

If GPR 29 does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Operation:


```
PC ← GPR[31]GPRLEN-1..1 || 0
if ( Config3ISA > 1 )
    ISAMode ← GPR[31]0
```

```
endif
if NotWordValue(GPR[29]) then
    UNPREDICTABLE
endif
temp ← GPR[29] + zero_extend(immediate || 02)
GPR[29] ← sign_extend(temp31..0)
```

Exceptions:

None.

Programming Notes:



Format: JRC16 rs

microMIPS

Purpose: Jump Register Compact (16-bit instr size)

To execute a branch to an instruction address in a register

Description: $PC \leftarrow GPR[rs]$

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors which implement MIPS64 and the ISAMode bit of the target address is MIPS64 (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS64 ISA, if the intended target ISAMode is MIPS64 (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump; that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:


This instruction has been recoded for Release 6.

Operation:

```
temp ← GPR[rs]
if Config3ISA = 1 then
    PC ← temp
else
    PC ← tempGPRELEN-1..1 || 0
    ISAMode ← temp0
endif
```

Exceptions:

None



Format: LBU16 rt, decoded_offset(base)

microMIPS

Purpose: Load Byte Unsigned (16-bit instr size)

To load a byte from memory as an unsigned value

Description: $GPR[rt] \leftarrow memory[GPR[base] + decoded_offset]$

The encoded offset field is decoded to get the actual offset value. This decoded value is added to the contents of base register to create the effective address. [Table 5.11](#) shows the encoded and decode values of the offset field.

Table 5.11 Offset Field Encoding Range -1, 0..14

Encoded Input (Hex)	Decoded Value (Decimal)
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
a	10
b	11
c	12
d	13
e	14
f	-1

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, zero-extended, and placed in GPR *rt*. The 4-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Operation:

$decoded_offset \leftarrow Decode(encoded_offset)$



```

vAddr ← sign_extend(decoded_offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
memdoubleword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor BigEndianCPU3
GPR[rt] ← zero_extend(memdoubleword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LDM {sregs, } {ra}, offset(base)

microMIPS

Purpose: Load Doubleword Multiple

To load a sequence of consecutive doublewords from memory

Description: {GPR[16], {GPR[17], {GPR[18], {GPR[19], {GPR[20], {GPR[21], {GPR[22], {GPR[23], {GPR[30]}}}}}}}}{GPR[31]} ← memory[GPR[base]+offset], ..., memory[GPR[base]+offset+8*(fn(reglist))]

The contents of consecutive 64-bit words at the memory location specified by the naturally aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in the GPRs defined by *reglist*. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (Binary)	List of Registers Loaded
0 0 0 0 1	GPR[16]
0 0 0 1 0	GPR[16], GPR[17]
0 0 0 1 1	GPR[16], GPR[17], GPR[18]
0 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19]
0 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20]
0 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21]
0 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22]
0 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23]
0 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30]
1 0 0 0 0	GPR[31]
1 0 0 0 1	GPR[16], GPR[31]
1 0 0 1 0	GPR[16], GPR[17], GPR[31]
1 0 0 1 1	GPR[16], GPR[17], GPR[18], GPR[31]
1 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]
1 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[31]
1 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[31]
1 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[31]
1 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[31]
1 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30], GPR[31]
All other combinations	Reserved

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the left-most register on the list and ending with the right-most register on the list. The effective address is incremented

for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instruction is **UNDEFINED**, if *base* is included in *reglist*. Reason for this is to allow restartability of the operation if an interrupt or exception has aborted the operation in the middle.

Pre-Release 6: The effective address must be naturally aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

if (Are64bitOperationsEnabled()) then
    vAddr ← sign_extend(offset) + GPR[base]
    for i←0 to fn(reglist)
        (pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
        memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
        GPR[gpr(reglist,i)] ← memdoubleword
        vAddr ← vAddr + 8
    endfor
else
    SignalException(ReservedInstruction)
endif

function fn(list)
    fn ← number of entries in list - 1
endfunction


```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Reserved Instruction, Watch

Programming Notes:

This instruction executes for a variable number of cycles and performs a variable number of loads from memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.



Format: LDP rd, offset(base)

microMIPS

Purpose: Load Doubleword Pair

To load two consecutive doublewords from memory

Description: $GPR[rd], GPR[rd+1] \leftarrow memory[GPR[base] + offset]$

The contents of the two consecutive 64-bit words at the memory location specified by the aligned effective address are fetched and placed in GPR *rd* and (*rd+1*). The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instructions is **UNDEFINED** if *rd* equals \$31.

The behavior of the instruction is **UNDEFINED**, if *base* and *rd* are the same. Reason for this is to allow restartability of the operation if an interrupt or exception has aborted the operation in the middle.

Pre-Release 6: The effective address must be naturally aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

if (Are64bitOperationsEnabled()) then
    vAddr ← sign_extend(offset) + GPR[base]
    (pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)←
    memdoubleword ← LoadMemory (CCA, DOUBLEWORD, pAddr, vAddr, DATA)←
    GPR[rd] ← memdoubleword
    vAddr ← sign_extend(offset) + GPR[base] + 8
    (pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)←
    memdoubleword ← LoadMemory (CCA, DOUBLEWORD, pAddr, vAddr, DATA)←
    GPR[rd+1]← memdoubleword
else
    SignalException(ReservedInstruction)
endif


```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Reserved Instruction, Watch

Programming Notes:

This instruction may execute for a variable number of cycles and performs two loads from memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.



Format: LHU16 rt, left_shifted_offset(base)

microMIPS

Purpose: Load Halfword Unsigned (16-bit instr size)

To load a halfword from memory as an unsigned value

Description: $GPR[rt] \leftarrow memory[GPR[base] + (offset \times 2)]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 4-bit unsigned *offset* is left shifted by one bit and then added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← zero_extend(offset || 0) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian2 || 0))
memdoubleword ← LoadMemory (CCA, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← zero_extend(memdoubleword15+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LI16 rd, decoded_immediate

microMIPS

Purpose: Load Immediate Word (16-bit instr size)

To load a 6-bit constant into a register.

Description: $GPR[rd] \leftarrow decoded_immediate$

The 7-bit encoded Immediate field is decoded to obtain the actual immediate value. [Table 5.12](#) shows the encoded values of the Immediate field and the actual immediate values.

Table 5.12 LI16 -1, 0..126 Immediate Field Encoding Range

Encoded Input (Hex)	Decoded Value (Decimal)
0	0
1	1
2	2
3	3
...	...
7e	126
7f	-1

The actual decoded immediate value is sign-extended and placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Operation:


```

decoded_immediate ← Decode(encoded_immediate)
temp ← sign_extend(decoded_immediate)
GPR[rd] ← temp63..0

```

Exceptions:

None



Format: LWP rd, offset(base)

microMIPS

Purpose: Load Word Pair

To load two consecutive words from memory

Description: $GPR[rd], GPR[rd+1] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the two consecutive 32-bit words at the memory location specified by the 32-bit aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rd* and (*rd+1*). The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instructions is **UNPREDICTABLE** if *rd* equals r31.

The behavior of the instruction is **UNPREDICTABLE**, if *base* and *rd* are the same. Reason for this is to allow restartability of the operation if an interrupt or exception has aborted the operation in the middle.

Pre-Release 6: The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU || 02)
GPR[rd] ← sign_extend(memdoubleword31+8*byte..8*byte)

vAddr ← sign_extend(offset) + GPR[base] + 4
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU || 02)
GPR[rd+1] ← sign_extend(memdoubleword31+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LW16 *rt*, left_shifted_offset(*base*)

microMIPS

Purpose: Load Word (16-bit instr size)

To load a word from memory as a signed value

Description: $GPR[rt] \leftarrow memory[GPR[base] + (offset \times 4)]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 4-bit unsigned *offset* is left shifted by two bits and then is added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← zero_extend(offset || 02) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LWM32 {srel6, } {ra}, offset(base)

microMIPS

Purpose: Load Word Multiple

To load a sequence of consecutive words from memory

Description: {GPR[16], {GPR[17], {GPR[18], {GPR[19], {GPR[20], {GPR[21], {GPR[22], {GPR[23], {GPR[30]}}}}}}}}{GPR[31]} ←
memory[GPR[base]+offset], ..., memory[GPR[base]+offset+4*(fn(reglist))]

The contents of consecutive 32-bit words at the memory location specified by the 32-bit aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in the GPRs defined by *reglist*. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (binary)	List of Registers Loaded
0 0 0 0 1	GPR[16]
0 0 0 1 0	GPR[16], GPR[17]
0 0 0 1 1	GPR[16], GPR[17], GPR[18]
0 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19]
0 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20]
0 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21]
0 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22]
0 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23]
0 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30]
1 0 0 0 0	GPR[31]
1 0 0 0 1	GPR[16], GPR[31]
1 0 0 1 0	GPR[16], GPR[17], GPR[31]
1 0 0 1 1	GPR[16], GPR[17], GPR[18], GPR[31]
1 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]
1 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[31]
1 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[31]
1 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[31]
1 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[31]
1 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30], GPR[31]
All other combinations	Reserved

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the

left-most register on the list and ending with the right-most register on the list. The effective address is incremented for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instruction is **UNPREDICTABLE**, if *base* is included in *reglist*. Reason for this is to allow restartability of the operation if an interrupt or exception has aborted the operation in the middle.

Pre-Release 6: The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
for i ← 0 to fn(reglist)
    (pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
    pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
    memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
    byte ← vAddr2..0 xor (BigEndianCPU || 02)
    GPR[gpr(reglist,i)] ← sign_extend(memdoubleword31+8*byte..8*byte)
    vAddr ← vAddr + 4
endfor

function fn(list)
    fn ← (number of entries in list) - 1
endfunction

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch

POOL16C 010001	reglist	offset	LWM16 0010
6	2	4	4

Format: LWM16 s0, {s1, {s2, {s3,}}} ra, left_shifted_offset(sp)

microMIPS

Purpose: Load Word Multiple (16-bit)

To load a sequence of consecutive words from memory

Description: $GPR[16], \{GPR[17], \{GPR[18], \{GPR[19], \}\}\} GPR[31] \leftarrow \text{memory}[GPR[29] + (\text{offset} \ll 2)], \dots, \text{memory}[GPR[19] + (\text{offset} \ll 2) + 4 * (\text{fn}(\text{reglist}))]$

The contents of consecutive 32-bit words at the memory location specified by the 32-bit aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in the GPRs defined by *reglist*. The 4-bit unsigned *offset* is first left shifted by two bits and then added to the contents of GPR *sp* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (binary)	List of Registers Loaded
0 0	GPR[16], GPR[31]
0 1	GPR[16], GPR[17], GPR[31]
1 0	GPR[16], GPR[17], GPR[18], GPR[31]
1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the left-most register on the list and ending with the right-most register on the list. The effective address is incremented for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← zero_extend(offset || 02) + GPR[sp]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
for i ← 0 to fn(reglist)
    (pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
    pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
    memdoubleword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
    byte ← vAddr2..0 xor (BigEndianCPU || 02)


```

```
        GPR[gpr(reglist,i)] ← sign_extend(memdoubleword31+8*byte..8*byte)
        vAddr ← vAddr + 4
    endfor

function fn(list)
    fn ← number of entries in list - 1
endfunction
```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LWGP
 LW16 rt, left_shifted_offset(gp)

microMIPS

Purpose: Load Word from Global Pointer (16-bit instr size)

To load a word from memory as a signed value

Description: $GPR[rt] \leftarrow memory[GPR[28] + (offset \times 4)]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 7-bit signed *offset* is left shifted by two bits and then added to the contents of GPR 28 to form the effective address.

Restrictions:

The 3-bit register field can only specify GPRs \$2-\$7, \$16, \$17.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.


Operation:

```

vAddr ← sign_extend(offset || 02) + GPR[28]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)
  
```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LWSP

LW16 rt, left_shifted_offset(sp)

microMIPS

Purpose: Load Word from Stack Pointer (16-bit instr size)

To load a word from memory as a signed value

Description: $GPR[rt] \leftarrow \text{memory}[GPR[29] + (\text{offset} \times 4)]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 5-bit signed *offset* is left shifted by two bits, zero-extended and then is added to the contents of GPR 29 to form the effective address.

Restrictions:

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:


```

vAddr ← zero_extend(offset || 02) + GPR[29]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: MOVE16 rd, rs

microMIPS

Purpose: Move Register (16-bit instr size)

To copy one GPR to another GPR.

Description: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}]$

The contents of GPR *rs* are placed into GPR *rd*.

Restrictions:


None

Operation:

$\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}]$

Exceptions:

None



Format: MOVEP *rd*, *re*, *rs*, *rt*

microMIPS

Purpose: Move a Pair of Registers

To copy two GPRs to another two GPRs.

Description: $GPR[rd] \leftarrow GPR[rs]; GPR[re] \leftarrow GPR[rt];$

The contents of GPR *rs* are placed into GPR *rd*. The contents of GPR *rt* are placed into GPR *re*.

The register numbers *rd* and *re* are determined by the encoded *enc_dest* field:

Table 5.13 Encoded and Decoded Values of the Enc_Dest Field

Encoded Value of Instr _{9..7} (Decimal)	Encoded Value of Instr _{9..7} (Hex)	Decoded Value of <i>rd</i> (Decimal)	Decoded Value of <i>re</i> (Decimal)
0	0x0	5	6
1	0x1	5	7
2	0x2	6	7
3	0x3	4	21
4	0x4	4	22
5	0x5	4	5
6	0x6	4	6
7	0x7	4	7

The register numbers *rs* and *rt* are determined by the encoded *enc_rs* and *enc_rt* fields:

Table 5.14 Encoded and Decoded Values of the Enc_rs and Enc_rt Fields

Encoded Value of Instr _{6..4} (or Instr _{3..1}) (Decimal)	Encoded Value of Instr _{6..4} (or Instr _{3..1}) (Hex)	Decoded Value of <i>rt</i> (or <i>rs</i>) (Decimal)	Symbolic Name (From ArchDefs.h)
0	0x0	0	zero
1	0x1	17	s1
2	0x2	2	v0
3	0x3	3	v1
4	0x4	16	s0
5	0x5	18	s2
6	0x6	19	s3
7	0x7	20	s4

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The destination register pair field, *enc_dest*, can only specify the register pairs defined in [Table 5.13](#).

The source register fields *enc_rs* and *enc_rt* can only specify GPRs 0,2-3,16-20.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:
$$\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}]; \text{GPR}[\text{re}] \leftarrow \text{GPR}[\text{rt}]$$
Exceptions:

None

Release 6

POOL16C 010001	rt	rs	NOT16 0000
6	3	3	4

Format: NOT16 *rt*, *rs*

microMIPS

Purpose: Invert (16-bit instr size)

To do a bitwise logical inversion.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \text{ XOR } 0\text{xffffffffffffffff}$ Invert the contents of GPR *rs* in a bitwise fashion and place the result into GPR *rt*.**Restrictions:**

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \text{ xor } 0\text{xffffffffffffffff}$ **Exceptions:**

None

POOL16C 010001	rt	rs	OR16 1001
6	3	3	4

Format: OR16 *rt*, *rs*

microMIPS32

Purpose: Or (16-bit instr size)

To do a bitwise logical OR

Description: $GPR[rt] \leftarrow GPR[rs] \text{ or } GPR[rt]$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical OR operation. The result is placed into GPR *rt*.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Availability and Compatibility:


This instruction has been recoded for Release 6.

Operation:

$GPR[rt] \leftarrow GPR[rs] \text{ or } GPR[rt]$

Exceptions:

None



Format: SB16 rt, offset(base)

microMIPS

Purpose: Store Byte (16-bit instr size)

To store a byte to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 8-bit byte of GPR *rt* is stored in memory at the location specified by the effective address. The 4-bit unsigned *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit *base* register field can only specify GPRs \$2-\$7, \$16, \$17.

The 3-bit *rt* register field can only specify GPRs \$0, \$2-\$7, \$17.

Operation:


```

vAddr ← zero_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
bytesel ← vAddr_2..0 xor BigEndianCPU3
datadoubleword ← GPR[rt]_63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, BYTE, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Bus Error, Address Error, Watch



Format: SDBBP16 code

EJTAG+microMIPS

Purpose: Software Debug Breakpoint (16-bit instr size)

To cause a debug breakpoint exception

Description:

This instruction causes a debug exception, passing control to the debug exception handler. If the processor is executing in Debug Mode when the SDBBP instruction is executed, the exception is a Debug Mode Exception, which sets the Debug_{DExcCode} field to the value 0x9 (Bp). The code field can be used for passing information to the debug exception handler, and is retrieved by the debug exception handler only by loading the contents of the memory word containing the instruction, using the DEPC register. The CODE field is not used in any way by the hardware.

Restrictions:

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:


```

and Config5.SBRI = 0 If DebugDM = 0 then
    SignalDebugBreakpointException()
else
    SignalDebugModeBreakpointException()
endif

```

Exceptions:

Debug Breakpoint Exception
Debug Mode Breakpoint Exception



Format: SDM {sregs, } {ra}, offset(base)

microMIPS

Purpose: Store Doubleword Multiple

To store a sequence of consecutive doublewords to memory

Description: memory[GPR[base]+offset], ..., memory[GPR[base]+offset+4*(fn(reglist))] ← {GPR[16], {GPR[17], {GPR[18], {GPR[19], {GPR[20], {GPR[21], {GPR[22], {GPR[23], {GPR[30]}}}}}}}}{GPR[21]}

The contents of the 64-bit doublewords of the GPRs defined by *reglist* are stored in memory at the location specified by the aligned effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (binary)	List of Registers Stored
0 0 0 0 1	GPR[16]
0 0 0 1 0	GPR[16], GPR[17]
0 0 0 1 1	GPR[16], GPR[17], GPR[18]
0 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19]
0 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20]
0 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21]
0 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22]
0 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23]
0 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30]
1 0 0 0 0	GPR[31]
1 0 0 0 1	GPR[16], GPR[31]
1 0 0 1 0	GPR[16], GPR[17], GPR[31]
1 0 0 1 1	GPR[16], GPR[17], GPR[18], GPR[31]
1 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]
1 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[31]
1 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[31]
1 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[31]
1 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[31]
1 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30], GPR[31]
All other combinations	Reserved

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the left-most register on the list and ending with the right-most register on the list. The effective address is incremented

for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

Pre-Release 6: The effective address must be naturally aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

if Are64bitOperationsEnabled() then
    vAddr ← sign_extend(offset) + GPR[base]
    for i←0 to fn(reglist)
        (pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
        datadoubleword ← GPR[gpr(reglist,i)]
        StoreMemory (CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)
        vAddr ← vAddr + 8
    endfor
else
    SignalException(ReservedInstruction)
endif

function fn(list)
    fn ← number of entries in list - 1
endfunction


```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Reserved Instruction, Watch

Programming Notes:

This instruction executes for a variable number of cycles and performs a variable number of stores to memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.



Format: SDP rd, offset(base)

microMIPS

Purpose: Store Doubleword Pair

To store two consecutive doublewords to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rd}], \text{GPR}[\text{rd}+1]$

The contents of the 64-bit doublewords of GPR *rd* and GPR *rd+1* are stored in memory at the location specified by the aligned effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instructions is **UNDEFINED** if *rd* equals \$31.

Pre-Release 6: The effective address must be naturally aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

if Are64bitOperationsEnabled() then
    vAddr ← sign_extend(offset) + GPR[base]
    (pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)←←
    datadoubleword ← GPR[rd]
    StoreMemory(CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)
    vAddr ← sign_extend(offset) + GPR[base] + 8
    (pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)←←
    datadoubleword ← GPR[rd+1]
    StoreMemory(CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)
else
    SignalException(ReservedInstruction)
endif


```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Reserved Instruction, Watch

Programming Notes:

This instruction may execute for a variable number of cycles and may perform a variable number of stores to memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.



Format: SH16 rt, left_shifted_offset(base)

microMIPS

Purpose: Store Halfword (16-bit instr size)

To store a halfword to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + (\text{offset} \times 2)] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 16-bit halfword of register *rt* is stored in memory at the location specified by the aligned effective address. The 4-bit unsigned *offset* is left shifted by one bit and then added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit *base* register field can only specify GPRs \$2-\$7, \$16, \$17.

The 3-bit *rt* register field can only specify GPRs \$0, \$2-\$7, \$17.

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← zero_extend(offset || 0) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian2 || 0))
bytesel ← vAddr_2..0 xor (BigEndianCPU2 || 0)
datadoubleword ← GPR[rt]_63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, HALFWORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SLL16 rd, rt, decoded_sa

microMIPS

Purpose: Shift Word Left Logical (16-bit instr size)

To left-shift a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \ll decoded_sa$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeros into the emptied bits; the word result is sign-extended and placed in GPR *rd*. The bit-shift amount is specified by decoding the *encoded_sa* field. [Table 5.15](#) lists the encoded values of the *encoded_sa* field and the actual bit shift amount values.

Table 5.15 Shift Amount Field Encoding

Encoded Input (Hex)	Decoded Value (Decimal)
0	8
1	1
2	2
3	3
4	4
5	5
6	6
7	7

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Operation:

```

decoded_sa ← DECODE(encoded_sa)
s ← decoded_sa
temp ← GPR[rt](31-s)..0
```


 || 0^s
GPR[rd] ← sign_extend(temp)

Exceptions:

None

Programming Notes:

Unlike nearly all other word operations, the SLL input operand does not have to be a properly sign-extended word value to produce a valid sign-extended 32-bit result. The result word is always sign-extended into a 64-bit destination register; this instruction with a zero shift amount truncates a 64-bit value to 32 bits and sign-extends it.



Format: SRL16 rd, rt, decoded_sa

microMIPS

Purpose: Shift Word Right Logical (16-bit instr size)

To execute a logical right-shift of a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \gg decoded_sa$ (logical)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, inserting zeros into the emptied bits; the word result is sign-extended and placed in GPR *rd*. The bit-shift amount is specified by . by decoding the *encoded_sa* field. [Table 5.16](#) lists the encoded values of the *encoded_sa* field and the actual bit shift amount values.

Table 5.16 Shift Amount Field Encoding

Encoded Input (Hex)	Decoded Value (Decimal)
0	8
1	1
2	2
3	3
4	4
5	5
6	6
7	7

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
decoded_sa ← DECODE(encoded_sa)
s ← decoded_sa
temp ← 0s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)

```

Exceptions:

None



Format: SUBU16 rd, rs, rt

microMIPS

Purpose: Subtract Unsigned Word (16-bit instr size)

To subtract 32-bit integers

Description: $GPR[rd] \leftarrow GPR[rs] - GPR[rt]$

The 32-bit word value in GPR *rt* is subtracted from the 32-bit value in GPR *rs* and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No integer overflow exception occurs under any circumstances.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

On 64-bit processors, if either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← GPR[rs] - GPR[rt]
GPR[rd] ← sign_extend(temp)


```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. It is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: SW16 rt, left_shifted_offset(base)

microMIPS

Purpose: Store Word (16-bit instr size)

To store a word to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + (\text{offset} \times 4)] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 32-bit word of GPR *rt* is stored in memory at the location specified by the aligned effective address. The 4-bit unsigned *offset* is left-shifted by two bits and then added to the contents of GPR *base* to form the effective address.

Restrictions:

The 3-bit *base* register field can only specify GPRs \$2-\$7, \$16, \$17.

The 3-bit *rt* register field can only specify GPRs \$0, \$2-\$7, \$17.

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← zero_extend(offset || 02) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWSP rt, left_shifted_offset(base)

microMIPS

Purpose: Store Word to Stack Pointer (16-bit instr size)

To store a word to memory

Description: $\text{memory}[\text{GPR}[29] + (\text{offset} \times 4)] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 32-bit word of GPR *rt* is stored in memory at the location specified by the aligned effective address. The 5-bit signed *offset* is left shifted by two bits, zero-extended and then is added to the contents of GPR 29 to form the effective address.

Restrictions:

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:

```

vAddr ← zero_extend(offset || 02) + GPR[29]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

POOL16C 010001	reglist	offset	SWM16 1010
6	2	4	4

Format: SWM16 s0, {s1, {s2, {s3,}}} ra, left_shifted_offset(sp)

microMIPS

Purpose: Store Word Multiple (16-bit)

To store a sequence of consecutive words to memory

Description: $\text{memory}[\text{GPR}[29]], \dots, \text{memory}[\text{GPR}[29] + (\text{offset} \ll 2) + 4 * (2 + \text{fn}(\text{reglist}))] \leftarrow \text{GPR}[16], \{\text{GPR}[17], \{\text{GPR}[18], \{\text{GPR}[19], \}\}\} \text{GPR}[31]$

The least-significant 32-bit words of the GPRs defined by *reglist* are stored in memory at the location specified by the aligned effective address. The 4-bit unsigned *offset* is added to the contents of GPR *sp* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (binary)	List of Registers Stored
0 0	GPR[16], GPR[31]
0 1	GPR[16], GPR[17], GPR[31]
1 0	GPR[16], GPR[17], GPR[18], GPR[31]
1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the left-most register on the list and ending with the right-most register on the list. The effective address is incremented for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← zero_extend(offset || 02) + GPR[sp]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
for i ← 0 to fn(reglist)
    (pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
    pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
    bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
    datadoubleword ← GPR[gpr(reglist, i)]63-8*bytesel..0 || 08*bytesel
    StoreMemory(CCA, WORD, datadoubleword, pAddr, vAddr, DATA)
endfor


```

```
        vAddr ← vAddr + 4
    endfor

    function fn(list)
        fn ← number of entries in list - 1
    endfunction
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWM32 {sregs, } {ra}, offset(base)

microMIPS

Purpose: Store Word Multiple

To store a sequence of consecutive words to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}], \dots, \text{memory}[\text{GPR}[\text{base}] + \text{offset} + 4 * (\text{fn}(\text{reglist}))] \leftarrow \{ \text{GPR}[16], \{ \text{GPR}[17], \{ \text{GPR}[18], \{ \text{GPR}[19], \{ \text{GPR}[20], \{ \text{GPR}[21], \{ \text{GPR}[22], \{ \text{GPR}[23], \{ \text{GPR}[30] \} \} \} \} \} \} \{ \text{GPR}[31] \} \}$

The least-significant 32-bit words of the GPRs defined by *reglist* are stored in memory at the location specified by the aligned effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The following table shows the encoding of the *reglist* field.

<i>reglist</i> Encoding (binary)	List of Registers Loaded
0 0 0 0 1	GPR[16]
0 0 0 1 0	GPR[16], GPR[17]
0 0 0 1 1	GPR[16], GPR[17], GPR[18]
0 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19]
0 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20]
0 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21]
0 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22]
0 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23]
0 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30]
1 0 0 0 0	GPR[31]
1 0 0 0 1	GPR[16], GPR[31]
1 0 0 1 0	GPR[16], GPR[17], GPR[31]
1 0 0 1 1	GPR[16], GPR[17], GPR[18], GPR[31]
1 0 1 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[31]
1 0 1 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[31]
1 0 1 1 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[31]
1 0 1 1 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[31]
1 1 0 0 0	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[31]
1 1 0 0 1	GPR[16], GPR[17], GPR[18], GPR[19], GPR[20], GPR[21], GPR[22], GPR[23], GPR[30], GPR[31]
All other combinations	Reserved

The register numbers and the effective addresses are correlated using the order listed in the table, starting with the left-most register on the list and ending with the right-most register on the list. The effective address is incremented

for each subsequent register on the list.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

Pre-Release 6: The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
for i ← 0 to fn(reglist)
    (pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
    pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
    bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
    datadoubleword ← GPR[gpr(reglist,i)]63-8*bytesel..0 || 08*bytesel
    StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)
    vAddr ← vAddr + 4
endfor

function fn(list)
    fn ← (number of entries in list) - 1
endfunction

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWP rs1, offset(base)

microMIPS

Purpose: Store Word Pair

To store two consecutive words to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rs1}], \text{GPR}[\text{rs1}+1]$

The least-significant 32-bit words of GPR *rs1* and GPR *rs1+1* are stored in memory at the location specified by the aligned effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

It is implementation-specific whether interrupts are disabled during the sequence of operations generated by this instruction.

Restrictions:

The behavior of the instructions is **UNDEFINED** if *rd* equals \$31.

Pre-Release 6: The effective address must be 32-bit aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
bytesel ← vAddr_2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rs1]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

vAddr ← sign_extend(offset) + GPR[base] + 4
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
bytesel ← vAddr_2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rs1+1]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

POOL16C 010001	rt	rs	XOR16 1000
6	3	3	4

Format: XOR16 *rt*, *rs*

microMIPS

Purpose: Exclusive OR (16-bit instr size)

To do a bitwise logical Exclusive OR

Description: $GPR[rt] \leftarrow GPR[rs] \text{ XOR } GPR[rt]$

Combine the contents of GPR *rs* and GPR *rt* in a bitwise logical Exclusive OR operation and place the result into GPR *rt*.

Restrictions:

The 3-bit register fields can only specify GPRs \$2-\$7, \$16, \$17.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

$GPR[rt] \leftarrow GPR[rs] \text{ xor } GPR[rt]$


Exceptions:

None

5.5 Recoded MIPS Instructions

This section describes recoded 32-bit instructions from MIPS32 and MIPS64 instruction sets specifically for use as part of the microMIPS instruction set.

Only instructions supported in Release 6 are provided. [Section 2.7, "Release 6 of the MIPS Architecture,"](#) lists instructions that have been added, removed and recoded in Release 6.



Format: ABS.fmt
 ABS.S ft, fs
 ABS.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Absolute Value

Description: $FPR[ft] \leftarrow \text{abs}(FPR[fs])$

The absolute value of the value in FPR *fs* is placed in FPR *ft*. The operand and result are values in format *fmt*.

The *Cause* bits are ORed into the *Flag* bits if no exception is taken.

If $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$ then this operation is arithmetic. For this case, any NaN operand signals invalid operation.

If $FCSR_{ABS2008}=1$ then this operation is non-arithmetic. For this case, both regular floating point numbers and NaN values are treated alike, only the sign bit is affected by this instruction. No IEEE exception can be generated for this case.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Availability and Compatibility:

ABS.PS has been removed in Release 6.

Operation:


`StoreFPR(ft, fmt, AbsoluteValue(ValueFPR(fs, fmt)))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation



Format: ADD *rd*, *rs*, *rt*

microMIPS

Purpose: Add Word

To add 32-bit integers. If an overflow occurs, then trap.

Description: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs* to produce a 32-bit result.

- If the addition results in 32-bit 2's complement arithmetic overflow, the destination register is not modified and an Integer Overflow exception occurs.
- If the addition does not overflow, the 32-bit result is sign-extended and placed into GPR *rd*.

Restrictions:

If either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits _{63..31} equal), then the result of the operation is UNPREDICTABLE.

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← (GPR[rs]31 | GPR[rs]31..0) + (GPR[rt]31 | GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← sign_extend(temp31..0)
endif


```

Exceptions:

Integer Overflow

Programming Notes:

ADDU performs the same arithmetic operation but does not trap on overflow.



Format: ADD.fmt
 ADD.S fd, fs, ft
 ADD.D fd, fs, ft

microMIPS
 microMIPS

Purpose: Floating Point Add

To add floating point values.

Description: $FPR[fd] \leftarrow FPR[fs] + FPR[ft]$

The value in $FPR.ft$ is added to the value in $FPR.fs$. The result is calculated to infinite precision, rounded by using to the current rounding mode in *FCSR*, and placed into $FPR.fd$. The operands and result are values in format *fmt*.

The *Cause* bits are ORed into the *Flag* bits if no exception is taken.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operands must be values in format *fmt*. If the fields are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

Availability and Compatibility:

ADD.PS has been removed in Release 6.

Operation:


StoreFPR (fd, fmt, ValueFPR(fs, fmt) +_{fmt} ValueFPR(ft, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Inexact, Overflow, Underflow



Format: ADDIU *rt*, *rs*, *immediate*

microMIPS

Purpose: Add Immediate Unsigned Word

To add a constant to a 32-bit integer.

Description: $GPR[rt] \leftarrow GPR[rs] + \text{immediate}$

The 16-bit signed *immediate* is added to the 32-bit value in GPR *rs* and the 32-bit arithmetic result is sign-extended and placed into GPR *rt*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← GPR[rs] + sign_extend(immediate)
GPR[rt] ← sign_extend(temp31..0)


```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDIUPC rt,immediate

microMIPS32 Release 6

Purpose: Add Immediate to PC (unsigned - non-trapping)

Description: $GPR[rt] \leftarrow (PC \& \sim 0x3 + \text{sign_extend}(\text{immediate} \ll 2))$

This instruction performs a PC-relative address calculation. The 19-bit immediate is shifted left by 2 bits, sign-extended, and added to the address of the ADDIUPC instruction. The result is placed in GPR *rt*.

This instruction is both a 32-bit and a 64-bit instruction. The 64-bit result is sign-extended by the same rules that govern sign-extension of virtual addresses in the MIPS64 Architecture, as described by the function `effective_address()` in the Privileged Resource Architecture.

Restrictions:

None

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Operation:


$GPR[rst] \leftarrow (PC \& \sim 0x3 + \text{sign_extend}(\text{immediate} \ll 2))$

Exceptions:

None

Programming Notes:

The term “unsigned” in this instruction mnemonic is a misnomer. “Unsigned” here means “non-trapping”. It does not trap on a signed 32-bit overflow. ADDIUPC corresponds to unsigned ADDIU, which does not trap on overflow, as opposed to ADDI, which does trap on overflow.



Format: ADDU rd, rs, rt

microMIPS

Purpose: Add Unsigned Word

To add 32-bit integers.

Description: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs* and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

If either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```


if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp  $\leftarrow$  GPR[rs] + GPR[rt]
GPR[rd]  $\leftarrow$  sign_extend(temp31..0)
    
```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ALIGN DALIGN
ALIGN rd,rs,rt,bp
DALIGN rd,rs,rt,bp

microMIPS32 Release 6
microMIPS64 Release 6

Purpose: Concatenate two GPRs, and extract a contiguous subset at a byte position

Description: $GPR[rd] \leftarrow (GPR[rt] \ll (8 * bp)) \text{ or } (GPR[rs] \gg (GPRLEN - 8 * bp))$

The input registers GPR *rt* and GPR *rs* are concatenated, and a register width contiguous subset is extracted, which is specified by the byte pointer *bp*.

The ALIGN instruction operates on 32-bit words, and has a 2-bit byte position field *bp*.

The DALIGN instruction operates on 64-bit doublewords, and has a 3-bit byte position field *bp*.

- ALIGN: The rightmost 32-bit word in GPR *rt* is left shifted as a 32-bit value by *bp* byte positions. The rightmost 32-bit word in register *rs* is right shifted as a 32-bit value by $(4 - bp)$ byte positions. These shifts are logical shifts, zero-filling. The shifted values are then *or*-ed together to create a 32-bit result that is sign-extended to 64-bits and written to destination GPR *rd*.
- DALIGN: The 64-bit doubleword in GPR *rt* is left shifted as a 64 bit value by *bp* byte positions. The 64-bit word in register *rs* is right shifted as a 64-bit value by $(8 - bp)$ byte positions. These shifts are logical shifts, zero-filling. The shifted values are then *or*-ed together to create a 64-bit result and written to destination GPR *rd*.

Restrictions:

Executing ALIGN and DALIGN with shift count *bp*=0 acts like a register to register move operation, and is redundant, and therefore discouraged. Software should not generate ALIGN or DALIGN with shift count *bp*=0.

DALIGN: A Reserved Instruction exception is signaled if access to 64-bit operations is not enabled.

Availability and Compatibility:

The ALIGN instruction is introduced by and required as of Release 6.

The DALIGN instruction is introduced by and required as of Release 6.

Programming Notes:

Release 6 ALIGN instruction corresponds to the pre-Release 6 DSP Module BALIGN instruction, except that BALIGN with shift counts of 0 and 2 are specified as being UNPREDICTABLE, whereas ALIGN (and DALIGN) defines all *bp* values, discouraging only *bp*=0.

Graphically,

Figure 5.1 ALIGN operation (32-bit)





Figure 5.2 DALIGN operation (64-bit)

**Operation:**

ALIGN:

```
tmp_rt_hi ← unsigned_word(GPR[rt]) << (8*bp)
tmp_rs_lo ← unsigned_word(GPR[rs]) >> (8*(4-bp))
tmp ← tmp_rt_hi or tmp_rs_lo
```

ALIGN on a 32-bit CPU:

```
GPR[rd] ← tmp
```

ALIGN on a 64-bit CPU:

```
GPR[rd] ← sign_extend.32(tmp)
```


DALIGN:

```
if not Are64bitOperationsEnabled()
  then SignalException(ReservedInstruction) endif
tmp_rt_hi ← unsigned_doubleword(GPR[rt]) << (8*bp)
tmp_rs_lo ← unsigned_doubleword(GPR[rs]) >> (8*(8-bp))
tmp ← tmp_rt_hi or tmp_rs_lo
GPR[rd] ← tmp
/* end of instruction */
```

Exceptions:

ALIGN: None

DALIGN: Reserved Instruction



Format: ALUIPC *rt*, *immediate*

microMIPS32 Release 6

Purpose: Aligned Add Upper Immediate to PC

Description: $GPR[rt] \leftarrow \sim 0x0FFFF \& (PC + \text{sign_extend}(immediate \ll 16))$

This instruction performs a PC-relative address calculation. The 16-bit immediate is shifted left by 16 bits, sign-extended, and added to the address of the ALUIPC instruction. The low 16 bits of the result are cleared, that is the result is aligned on a 64K boundary. The result is placed in GPR *rt*.

This instruction is both a 32-bit and a 64-bit instruction. The 64-bit result is sign-extended by the same rules that govern sign-extension of virtual addresses in the MIPS64 Architecture, as described by the function `effective_address()` in the Privileged Resource Architecture.

Restrictions:

None

Availability and Compatibility:


This instruction is introduced by and required as of Release 6.

Operation:

$GPR[rt] \leftarrow \sim 0x0FFFF \& (PC + \text{sign_extend}(immediate \ll 16))$

Exceptions:

None



Format: AND rd, rs, rt

microMIPS

Purpose: and

To do a bitwise logical AND.

Description: $GPR[rd] \leftarrow GPR[rs] \text{ and } GPR[rt]$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical AND operation. The result is placed into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ and } GPR[rt]$

Exceptions:

None



Format: ANDI *rt*, *rs*, *immediate*

microMIPS

Purpose: and immediate

To do a bitwise logical AND with a constant

Description: $GPR[rt] \leftarrow GPR[rs] \text{ and } \text{zero_extend}(\text{immediate})$

The 16-bit immediate is zero-extended to the left and combined with the contents of GPR *rs* in a bitwise logical AND operation. The result is placed into GPR *rt*.

Restrictions:

None

Operation:

$GPR[rt] \leftarrow GPR[rs] \text{ and } \text{zero_extend}(\text{immediate})$

Exceptions:

None

31	26	25	21	20	16	15	0
AUI 000100	rt		rs		immediate		
DAUI 111100	rt		rs		immediate		
POOL32I 010000	DAHI 10001		rs		immediate		
POOL32I 010000	DATI 10000		rs		immediate		
6	5		5		16		

Format: AUI DAUI DAHI DATI

AUI rt, rs immediate

microMIPS32 Release 6

DAUI rt, rs immediate

microMIPS64 Release 6

DAHI rs, rs immediate

microMIPS64 Release 6

DATI rs, rs immediate

microMIPS64 Release 6

Purpose: Add Immediate to Upper Bits

AUI: Add Upper Immediate

DAUI: Doubleword Add Upper Immediate

DAHI: Doubleword Add Higher Immediate

DATI: Doubleword Add Top Immediate

Description:

```

AUI:  GPR[rt] ← sign_extend.32( GPR[rs] + sign_extend(immediate << 16) )
DAUI: GPR[rt] ← GPR[rs] + sign_extend(immediate << 16)
DAHI: GPR[rs] ← GPR[rs] + sign_extend(immediate << 32)
DATI: GPR[rs] ← GPR[rs] + sign_extend(immediate << 48)

```

AUI: The 16 bit immediate is shifted left 16 bits, sign-extended, and added to the register *rs*, storing the result in *rt*. AUI is a 32-bit compatible instruction, so on a 64-bit CPU the result is sign extended as if a 32-bits signed address.

DAHI: The 16 bit immediate is shifted left 32 bits, sign-extended, and added to the register *rs*, overwriting *rs* with the result.

DATI: The 16 bit immediate is shifted left 48bits, sign-extended, and added to the register *rs*, overwriting *rs* with the result.

DAUI: The 16-bit immediate is shifted left 16 bits, sign-extended, and added to the register *rs*; the results are stored in *rt*.

In Release 6, LUI is an assembly idiom for AUI with *rs*=0.

Restrictions:

DAUI: *rs* cannot be *r0*, the zero register. The encoding may be used for other instructions or must signal a Reserved

Instruction exception.

DAUI, DAHI, DATI: Reserved Instruction exception if 64-bit instructions are not enabled.

Availability and Compatibility:

AUI is introduced by and required as of Release 6.

DAUI is introduced by and required as of MIPS64 Release 6.

DAHI is introduced by and required as of MIPS64 Release 6.

DATI is introduced by and required as of MIPS64 Release 6.

Operation:

```
AUI:  GPR[rt] ← sign_extend.32( GPR[rs] + sign_extend(immediate << 16) )
DAUI: GPR[rt] ← GPR[rs] + sign_extend(immediate << 16)
DAHI: GPR[rs] ← GPR[rs] + sign_extend(immediate << 32)
DATI: GPR[rs] ← GPR[rs] + sign_extend(immediate << 48)
```

Exceptions:

AUI: None.

DAUI, DAHI, DATI: Reserved Instruction

Programming Notes:

AUI (and DAUI, DAHI and DATI on MIPS64 Release 6) can be used to synthesize large constants in situations where it is not convenient to load a large constant from memory. To simplify hardware that may recognize sequences of instructions as generating large constants, AUI/DAUI/DAHI/DATI should be used in a stylized manner.

To create an integer:

```
LUI rd, imm_low(rtmp)
ORI rd, rd, imm_upper
DAHI rd, imm_high
DATI rd, imm_top
```

To create a large offset for a memory access whose address is of the form `rbase+large_offset`:

```
AUI rtmp, rbase, imm_upper
DAHI rtmp, imm_high
DATI rtmp, imm_top
LW rd, (rtmp)imm_low
```


To create a large constant operand for an instruction of the form `rd:=rs+large_immediate`
or `rd:=rs-large_immediate`:

32-bits:

```
AUI  rtmp, rs,  imm_upper
ADDIU rd,  rtmp, imm_low
```

64-bits:

```
AUI  rtmp, rs, imm_upper
DAHI rtmp, imm_high
DATI rtmp, imm_top
DADDUI rd, rtmp, imm_low
```



Format: AUIPC rt, immediate

microMIPS32 Release 6

Purpose: Add Upper Immediate to PC

Description: $GPR[rt] \leftarrow (PC + \text{sign_extend}(\text{immediate} \ll 16))$

This instruction performs a PC-relative address calculation. The 16-bit immediate is shifted left by 16 bits, sign-extended, and added to the address of the AUIPC instruction. The result is placed in GPR *rt*.

In a MIPS64 implementation, the 32-bit result is sign extended from bit 31 to bit 63.

This instruction is both a 32-bit and a 64-bit instruction. The 64-bit result is sign-extended by the same rules that govern sign-extension of virtual address in the MIPS64 Architecture, as described by the function `effective_address()` in the Privileged Resource Architecture.

Restrictions:

None

Availability and Compatibility:


This instruction is introduced by and required as of Release 6.

Operation:

$GPR[rt] \leftarrow (PC + \text{sign_extend}(\text{immediate} \ll 16))$

Exceptions:

None



Format: BALC offset

microMIPS32 Release 6

Purpose: Branch and Link, Compact

To do an unconditional PC-relative procedure call.

Description: `procedure_call` (no delay slot)

Place the return address link in GPR 31. The return link is the address of the instruction immediately following the branch, where execution continues after a procedure call. (Because compact branches have no delay slots, see below.)

A 27-bit signed offset (the 26-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed when the branch is taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Exceptions:

None

Operation:

```
target_offset ← sign_extend( offset || 01 )
GPR[31] ← PC+4
PC ← PC+4 + sign_extend(target_offset)
```

31	26	25	21	20	16	15	0
POOL32I 010001	BC1EQZC 01000		ft				offset
POOL32I 010001	BC1NEZC 01001		ft				offset
6	5		5				16

Format: BC1EQZC BC1NEZC
 BC1EQZC ft, offset
 BC1NEZC ft, offset

microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Branch if Coprocessor 1 (FPU) Register Bit 0 Equal/Not Equal to Zero

BC1EQZC: Branch if Coprocessor 1 (FPU) Register Bit 0 is Equal to Zero

BC1NEZC: Branch if Coprocessor 1 (FPR) Register Bit 0 is Not Equal to Zero

Description:

BC1EQZC: if FPR[ft] & 1 = 0 then branch
 BC1NEZC: if FPR[ft] & 1 ≠ 0 then branch

The condition is evaluated on FPU register *ft*.

- For BC1EQZC, the condition is true if and only if bit 0 of the FPU register *ft* is zero.
- For BC1NEZC, the condition is true if and only if bit 0 of the FPU register *ft* is non-zero.

If the condition is false, the branch is not taken, and execution continues with the next instruction.

A 17-bit signed offset (the 16-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed if the branch is taken.

Restrictions:

If access to Coprocessor 1 is not enabled, a Coprocessor Unusable Exception is signaled.

Because these instructions BC1EQZC and BC1NEZC do not depend on a particular floating point data type, they operate whenever Coprocessor 1 is enabled.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Exceptions:

Coprocessor Unusable¹

Operation:

$\text{tmp} \leftarrow \text{ValueFPR}(\text{ft}, \text{UNINTERPRETED_WORD})$

1. In Release 6, BC1EQZC and BC1NEZC are required, if the FPU is implemented. They must not signal a Reserved Instruction exception. They can signal a Coprocessor Unusable Exception.

```
BC1EQZC: cond ← tmp & 1 = 0
BC1NEZC: cond ← tmp & 1 ≠ 0
if cond then
    target_PC ← ( PC+4 + sign_extend( offset << 1 ) )
    PC ← target_PC
```

Programming Notes:

Release 6: These instructions, BC1EQZC and BC1NEZC, replace the pre-Release 6 instructions BC1F and BC1T. These Release 6 FPU branches depend on bit 0 of the scalar FPU register.

Note: BC1EQZC and BC1NEZC do not have a format or data type width. The same instructions are used for branches based on conditions involving any format, including 32-bit S (single precision) and W (word) format, and 64-bit D (double precision) and L (longword) format, as well as 128-bit MSA. The FPU scalar comparison instructions *CMP.condn.fmt* produce an all ones or all zeros truth mask of their format width with the upper bits (where applicable) UNPREDICTABLE. BC1EQZ and BC1NEZ consume only bit 0 of the *CMP.condn.fmt* output value, and therefore operate correctly independent of *fmt*.

31	26	25	21	20	16	15	0
POOL32I	BC2EQZC 01010		ct		offset		
POOL32I	BC2NEZC 01011		ct		offset		
6	5		5		16		

Format: BC2EQZC BC2NEZC
 BC2EQZC ct, offset
 BC2NEZC ct, offset

microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Branch if Coprocessor 2 Condition (Register) Equal/Not Equal to Zero

BC2EQZC: Branch if Coprocessor 2 Condition (Register) is Equal to Zero

BC2NEZC: Branch if Coprocessor 2 Condition (Register) is Not Equal to Zero

Description:

BC2EQZC: if COP2Condition[ct] = 0 then branch
 BC2NEZC: if COP2Condition[ct] ≠ 0 then branch

The 5-bit field `ct` specifies a coprocessor 2 condition.

- For BC2EQZC if the coprocessor 2 condition is true the branch is taken.
- For BC2NEZC if the coprocessor 2 condition is false the branch is taken.

A 17-bit signed offset (the 16-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

Compact branches do not have delay slots. The instruction after the branch is NOT executed if the branch is taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

If access to Coprocessor 2 is not enabled, a Coprocessor Unusable Exception is signaled.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Exceptions:

Coprocessor Unusable, Reserved Instruction

Operation:

```

tmpcond ← Coprocessor2Condition(ct)
if BC2EQZC then
    tmpcond ← not(tmpcond)
endif

if tmpcond then
    PC ← PC+4 + sign_extend( immediate << 1 )
endif

```

Implementation Notes:

As of Release 6 these instructions, BC2EQZC and BC2NEZC, replace the pre-Release 6 instructions BC2F and BC2T, which had a 3-bit condition code field (as well as nullify and true/false bits). Release 6 makes all 5 bits of the `cc` condition code available to the coprocessor designer as a condition specifier.

A customer defined coprocessor instruction set can implement any sort of condition it wants. For example, it could implement up to 32 single-bit flags, specified by the 5-bit field `cc`. It could also implement conditions encoded as values in a coprocessor register (such as testing the least significant bit of a coprocessor register) as done by Release 6 instructions BC1EQZ/BC1NEZ.

31	26	25	21	20	16	15	0
POP60 110000	BLEZALC rt ≠ 00000 00000				offset		
POP60 110000	BGEZALC rs = rt ≠ 00000 rt rs				offset		
POP70 111000	BGTZALC rt ≠ 00000 00000				offset		
POP70 111000	BLTZALC rs = rt ≠ 00000 rt rs				offset		
POP35 011101	BEQZALC rs < rt rt ≠ 00000 00000				offset		
POP37 011111	BNEZALC rs < rt rt ≠ 00000 00000				offset		
6	5		5		16		

Format: B{LE,GE,GT,LT,EQ,NE}ZALC

BLEZALC rt, offset

BGEZALC rt, offset

BGTZALC rt, offset

BLTZALC rt, offset

BEQZALC rt, offset

BNEZALC rt, offset

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS32 Release 6

Purpose: Compact Zero-Compare and Branch-and-Link Instructions

BLEZALC: Compact branch-and-link if GPR rt is less than or equal to zero

BGEZALC: Compact branch-and-link if GPR rt is greater than or equal to zero

BGTZALC: Compact branch-and-link if GPR rt is greater than zero

BLTZALC: Compact branch-and-link if GPR rt is less than to zero

BEQZALC: Compact branch-and-link if GPR rt is equal to zero

BNEZALC: Compact branch-and-link if GPR rt is not equal to zero

Description: if condition(GPR[rt]) then procedure_call branch

The condition is evaluated. If the condition is true, the branch is taken.

Places the return address link in GPR 31. The return link is the address of the instruction immediately following the branch, where execution continues after a procedure call.

The return address link is unconditionally updated.

A 17-bit signed offset (the 16-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

BLEZALC: the condition is true if and only if GPR rt is less than or equal to zero.

BGEZALC: the condition is true if and only if GPR rt is greater than or equal to zero.

BLTZALC: the condition is true if and only if GPR rt is less than zero.

BGTZALC: the condition is true if and only if GPR rt is greater than zero.

BEQZALC: the condition is true if and only if GPR rt is equal to zero.

BNEZALC: the condition is true if and only if GPR rt is not equal to zero.

Compact branches do not have delay slots. The instruction after a compact branch is only executed if the branch is not taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Exceptions:

None

Operation:

```
GPR[31] ← PC+4
target_offset ← sign_extend( offset || 01 )

BLTZALC: cond ← GPR[rt] < 0
BLEZALC: cond ← GPR[rt] ≤ 0
BGEZALC: cond ← GPR[rt] ≥ 0
BGTZALC: cond ← GPR[rt] > 0
BEQZALC: cond ← GPR[rt] = 0
BNEZALC: cond ← GPR[rt] ≠ 0

if cond then
    PC ← ( PC+4+ sign_extend( target_offset ) )
endif
```

Programming Notes:

Software that performs incomplete instruction decode may incorrectly decode these new instructions, because of their very tight encoding. For example, a disassembler might look only at the primary opcode field, instruction bits 31-26, to decode BLEZL without checking that the “rt” field is zero. Such software violated the pre-Release 6 architecture specification.

With the 16-bit offset shifted left 2 bits and sign extended, the conditional branch range is ± 128 KBytes. Other instructions such as pre-Release 6 JAL and JALR, or Release 6 JIALC and BALC have larger ranges. In particular, BALC, with a 26-bit offset shifted by 2 bits, has a 28-bit range, ± 128 MBytes. Code sequences using AUIPC, DAHI, DATI, and JIALC allow still greater PC-relative range.

31	26	25	21	20	16	15	0
POP75 111101	BLEZC				offset		
	rt ≠ 00000		00000				
POP75 111101	BGEZC rs = rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP75 111101	BGEC (BLEC) rs ≠ rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP65 110101	BGTZC				offset		
	rt ≠ 00000		00000				
POP65 110101	BLTZC rs = rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP65 110101	BLTC (BGTC) rs ≠ rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP60 110000	BGEUC (BLEUC) rs ≠ rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP70 111000	BLTUC (BGTUC) rs ≠ rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP35 011101	BEQC rs < rt				offset		
	rt ≠ 00000		rs ≠ 00000				
POP37 011111	BNEC rs < rt				offset		
	rt ≠ 00000		rs ≠ 00000				
6	5		5		16		

31	26	25	21	20	0
POP40 100000	BEQZC		offset		
	rs ≠ 00000				
	rs				
POP50 101000	BNEZC		offset		
	rs ≠ 00000				
	rs				
6	5		21		

Format: B<cond>C rs, rt, offset

microMIPS32 Release 6

Purpose: Compact Compare-and-Branch Instructions

Format Details:

Equal/Not-Equal register-register compare and branch with 16-bit offset:

BEQC rs, rt, offset

BNEC rs, rt, offset

microMIPS32 Release 6

microMIPS32 Release 6

Signed register-register compare and branch with 16-bit offset:

BLTC *rs*, *rt*, *offset*
BGEC *rs*, *rt*, *offset*

microMIPS32 Release 6
microMIPS32 Release 6

Unsigned register-register compare and branch with 16-bit offset:

BLTUC *rs*, *rt*, *offset*
BGEUC *rs*, *rt*, *offset*

microMIPS32 Release 6
microMIPS32 Release 6

Assembly idioms with reversed operands for signed/unsigned compare-and-branch:

BGTC *rt*, *rs*, *offset*
BLEC *rt*, *rs*, *offset*
BGTUC *rt*, *rs*, *offset*
BLEUC *rt*, *rs*, *offset*

Assembly Idiom
Assembly Idiom
Assembly Idiom
Assembly Idiom

Signed Compare register to Zero and branch with 16-bit offset:

BLTZC *rt*, *offset*
BLEZC *rt*, *rs*, *offset*
BGEZC *rt*, *offset*
BGTZC *rt*, *rs*, *offset*

microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6

Equal/Not-equal Compare register to Zero and branch with 21-bit offset:

BEQZC *rt*, *rs*, *offset*
BNEZC *rt*, *rs*, *offset*

microMIPS32 Release 6
microMIPS32 Release 6

Description: if condition(GPR[*rs*] and/or GPR[*rt*]) then compact branch

The condition is evaluated. If the condition is true, the branch is taken.

An 18/23-bit signed offset (the 16/21-bit offset field shifted left 2 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

The offset is 16 bits for most compact branches, including BLTC, BLEC, BGEC, BGTC, BNEQC, BNEC, BLTUC, BLEUC, BGEUC, BGTC, BLTZC, BLEZC, BGEZC, BGTZC. The offset is 21 bits for BEQZC and BNEZC.

Compact branches have no delay slot: the instruction after the branch is NOT executed if the branch is taken.

The conditions are as follows:

Equal/Not-equal register-register compare-and-branch with 16-bit offset:

BEQC: Compact branch if GPRs are equal
BNEC: Compact branch if GPRs are not equal

Signed register-register compare and branch with 16-bit offset:

BLTC: Compact branch if GPR *rs* is less than GPR *rt*
BGEC: Compact branch if GPR *rs* is greater than or equal to GPR *rt*

Unsigned register-register compare and branch with 16-bit offset:

BLTUC: Compact branch if GPR *rs* is less than GPR *rt*, unsigned
BGEUC: Compact branch if GPR *rs* is greater than or equal to GPR *rt*, unsigned

Assembly Idioms with Operands Reversed:

BLEC: Compact branch if GPR *rt* is less than or equal to GPR *rs* (alias for BGEC)
BGTC: Compact branch if GPR *rt* is greater than GPR *rs* (alias for BLTC)
BLEUC: Compact branch if GPR *rt* is less than or equal to GPR *rt*, unsigned (alias for BGEUC)
BGTUC: Compact branch if GPR *rt* is greater than GPR *rs*, unsigned (alias for BLTUC)

Compare register to zero and branch with 16-bit offset:

BLTZC: Compact branch if GPR *rt* is less than zero
 BLEZC: Compact branch if GPR *rt* is less than or equal to zero
 BGEZC: Compact branch if GPR *rt* is greater than or equal to zero
 BGTZC: Compact branch if GPR *rt* is greater than zero

Compare register to zero and branch with 21-bit offset:

BEQZC: Compact branch if GPR *rs* is equal to zero
 BNEZC: Compact branch if GPR *rs* is not equal to zero

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Exceptions:

None

Operation:

```
target_offset ← sign_extend( offset || 01 )


/* Register-register compare and branch, 16 bit offset: */
/* Equal / Not-Equal */
BEQC: cond ← GPR[rs] = GPR[rt]
BNEC: cond ← GPR[rs] ≠ GPR[rt]
/* Signed */
BLTC: cond ← GPR[rs] < GPR[rt]
BGEC: cond ← GPR[rs] ≥ GPR[rt]
/* Unsigned: */
BLTUC: cond ← unsigned(GPR[rs]) < unsigned(GPR[rt])
BGEUC: cond ← unsigned(GPR[rs]) ≥ unsigned(GPR[rt])

/* Compare register to zero, small offset: */
BLTZC: cond ← GPR[rt] < 0
BLEZC: cond ← GPR[rt] ≤ 0
BGEZC: cond ← GPR[rt] ≥ 0
BGTZC: cond ← GPR[rt] > 0
/* Compare register to zero, large offset: */
BEQZC: cond ← GPR[rs] = 0
BNEZC: cond ← GPR[rs] ≠ 0

if cond then
  PC ← ( PC+4+ sign_extend( offset ) )
end if
```

Programming Notes:

Legacy software that performs incomplete instruction decode may incorrectly decode these new instructions, because of their very tight encoding. For example, a disassembler that looks only at the primary opcode field (instruction bits 31-26) to decode BLEZL without checking that the “rt” field is zero violates the pre-Release 6 architecture specification. Complete instruction decode allows reuse of pre-Release 6 BLEZL opcode for Release 6 conditional branches.



Format: BC offset

microMIPS32 Release 6

Purpose: Branch, Compact

Description: $PC \leftarrow PC+4 + \text{sign_extend}(\text{offset} \ll 1)$

A 27-bit signed offset (the 26-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

Compact branches have no delay slot: the instruction after the branch is NOT executed when the branch is taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:


This instruction is introduced by and required as of Release 6.

Exceptions:

None

Operation:

```
target_offset ← sign_extend( offset || 01 )
PC ← ( PC+4 + sign_extend(target_offset) )
```

Format: BREAK

microMIPS

Purpose: Breakpoint

To cause a Breakpoint exception

Description:

A breakpoint exception occurs, immediately and unconditionally transferring control to the exception handler. The *code* field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Restrictions:


None

Operation:

`SignalException(Breakpoint)`

Exceptions:

Breakpoint



Format: BITSWAP DBITSWAP
 BITSWAP rd,rt
 DBITSWAP rd,rt

microMIPS32 Release 6
microMIPS64 Release 6

Purpose: Swaps (reverses) bits in each byte

Description: $\text{GPR}[\text{rd}].\text{byte}(i) \leftarrow \text{reverse_bits_in_byte}(\text{GPR}[\text{rt}].\text{byte}(i))$, for all bytes i

Each byte in input GPR rt is moved to the same byte position in output GPR rd , with bits in each byte reversed.

BITSWAP is a 32-bit instruction. BITSWAP operates on all 4 bytes of a 32-bit GPR on a 32-bit CPU. On a 64-bit CPU, BITSWAP operates on the low 4 bytes, sign extending to 64-bits.

DBITSWAP operates on all 8 bytes of a 64-bit GPR on a 64-bit CPU.

Restrictions:

BITSWAP: None.

Availability and Compatibility:

The BITSWAP instruction is introduced by and required as of Release 6.

The DBITSWAP instruction is introduced by and required as of Release 6.

Operation:

```
BITSWAP:
  for i in 0 to 3 do /* for all bytes in 32-bit GPR width */
    tmp.byte(i) ← reverse_bits_in_byte( GPR[rt].byte(i) )
  endfor
  GPR[rd] ← sign_extend.32( tmp )
```

```
DBITSWAP:
  for i in 0 to 7 do /* for all bytes in 64-bit GPR width */
    tmp.byte(i) ← reverse_bits_in_byte( GPR[rt].byte(i) )
  endfor
  GPR[rd] ← tmp
```

where

```
function reverse_bits_in_byte(inbyte)
  outbyte7 ← inbyte0
  outbyte6 ← inbyte1
  outbyte5 ← inbyte2
  outbyte4 ← inbyte3
  outbyte3 ← inbyte4
```

```
    outbyte2 ← inbyte5  
    outbyte1 ← inbyte6  
    outbyte0 ← inbyte7  
    return outbyte  
end function
```

Exceptions:

BITSWAP: None

DBITSWAP: Reserved Instruction.

Programming Notes:

The Release 6 BITSWAP instruction corresponds to the DSP Module BITREV instruction, except that the latter bit-reverses the least-significant 16-bit halfword of the input register, zero extending the rest, while BITSWAP operates on 32-bits.

31	26	25	21	20	16	15	0
POP35 011101		BOVC rs >=rt			offset		
		rt	rs				
POP37 011111		BNVC rs>=rt			offset		
		rt	rs				
6		5		5		16	

Format: BOVC BNVC
 BOVC $rt, rs, offset$
 BNVC $rt, rs, offset$

microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Branch on Overflow, Compact; Branch on No Overflow, Compact

BOVC: Detect overflow for add (signed 32 bits) and branch if overflow.

BNVC: Detect overflow for add (signed 32 bits) and branch if no overflow.

Description: branch if/if-not NotWordValue($GPR[rs] + GPR[rt]$)

- BOVC performs a signed 32-bit addition of rs and rt . BOVC discards the sum, but detects signed 32-bit integer overflow of the sum (and the inputs, in MIPS64), and branches if such overflow is detected.
- BNVC performs a signed 32-bit addition of rs and rt . BNVC discards the sum, but detects signed 32-bit integer overflow of the sum (and the inputs, in MIPS64), and branches if such overflow is not detected.

A 17-bit signed offset (the 16-bit offset field shifted left 1 bits) is added to the address of the instruction following the branch (not the branch itself), to form a PC-relative effective target address.

On 64-bit processors, BOVC and BNVC detect signed 32-bit overflow on the input registers as well as the output. This checking is performed even if 64-bit operations are not enabled.

The special case with $rt=0$ (for example, $GPR[0]$) is allowed. On MIPS64, this checks that the input value of rs is a well-formed signed 32-bit integer: BOVC $rs, r0, offset$ branches if rs is not a 32-bit integer, and BNVC $rs, r0, offset$ branches if rs is a 32-bit integer.

The special case of $rs=0$ and $rt=0$ is allowed. BOVC never branches, while BNVC always branches.

Compact branches do not have delay slots. The instruction after the branch is NOT executed if the branch is taken.

Restrictions:

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Operation:

```
input_overflow ← NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])
```


```
temp1 ← sign_extend.32( GPR[rs]31..0 )
temp2 ← sign_extend.32( GPR[rt]31..0 )
tempd ← temp1 + temp2 // wider than 32-bit precision
sum_overflow ← (tempd32 ≠ tempd31)
```

```
BOVC: cond ← sum_overflow or input_overflow
BNVC: cond ← not( sum_overflow or input_overflow )
```

```
if cond then
    PC ← ( PC+4 + sign_extend( offset << 1 ) )
endif
```

Exceptions:

None



Format: CACHE op, offset(base)

microMIPS

Purpose: Perform Cache Operation

To perform the cache operation specified by op.


Description:

The 9-bit offset is sign-extended and added to the contents of the base register to form an effective address. The effective address is used in one of the following ways based on the operation to be performed and the type of cache as described in the following table.

Table 5.17 Usage of Effective Address

Operation Requires an	Type of Cache	Usage of Effective Address
Address	Virtual	The effective address is used to address the cache. An address translation may or may not be performed on the effective address (with the possibility that a TLB Refill or TLB Invalid exception might occur)
Address	Physical	The effective address is translated by the MMU to a physical address. The physical address is then used to address the cache
Index	N/A	<p>The effective address is translated by the MMU to a physical address. It is implementation dependent whether the effective address or the translated physical address is used to index the cache. As such, an unmapped address (such as within kseg0) should always be used for cache operations that require an index. See the Programming Notes section below.</p> <p>Assuming that the total cache size in bytes is CS, the associativity is A, and the number of bytes per tag is BPT, the following calculations give the fields of the address which specify the way and the index:</p> $\begin{aligned} \text{OffsetBit} &\leftarrow \text{Log2}(\text{BPT}) \\ \text{IndexBit} &\leftarrow \text{Log2}(\text{CS} / \text{A}) \\ \text{WayBit} &\leftarrow \text{IndexBit} + \text{Ceiling}(\text{Log2}(\text{A})) \\ \text{Way} &\leftarrow \text{Addr}_{\text{WayBit}-1..\text{IndexBit}} \\ \text{Index} &\leftarrow \text{Addr}_{\text{IndexBit}-1..\text{OffsetBit}} \end{aligned}$ <p>For a direct-mapped cache, the Way calculation is ignored and the Index value fully specifies the cache tag. This is shown symbolically in the figure below.</p>

Figure 5.3 Usage of Address Fields to Select Index and Way



A TLB Refill and TLB Invalid (both with cause code equal TLBL) exception can occur on any operation. For index

operations (where the address is used to index the cache but need not match the cache tag), software must use unmapped addresses to avoid TLB exceptions. This instruction never causes TLB Modified exceptions nor TLB Refill exceptions with a cause code of TLBS. This instruction never causes Execute-Inhibit nor Read-Inhibit exceptions.

The effective address may be an arbitrarily-aligned by address. The CACHE instruction never causes an Address Error Exception due to a non-aligned address.

As a result, a Cache Error exception may occur because of some operations performed by this instruction. For example, if a Writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Also, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error. However, cache error exceptions must not be triggered by an Index Load Tag or Index Store tag operation, as these operations are used for initialization and diagnostic purposes.

An Address Error Exception (with cause code equal AdEL) may occur if the effective address references a portion of the kernel address space which would normally result in such an exception. It is implementation dependent whether such an exception does occur.

It is implementation dependent whether a data watch is triggered by a cache instruction whose address matches the Watch register address match conditions.

The CACHE instruction and the memory transactions which are sourced by the CACHE instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

Bits [22:21] of the instruction specify the cache on which to perform the operation, as follows:

Table 5.18 Encoding of Bits[17:16] of CACHE Instruction

Code	Name	Cache
0b00	I	Primary Instruction
0b01	D	Primary Data or Unified Primary
0b10	T	Tertiary
0b11	S	Secondary

Bits [25:23] of the instruction specify the operation to perform. To provide software with a consistent base of cache operations, certain encodings must be supported on all processors. The remaining encodings are recommended

When implementing multiple level of caches and where the hardware maintains the smaller cache as a proper subset of a larger cache (every address which is resident in the smaller cache is also resident in the larger cache; also known as the inclusion property). It is recommended that the CACHE instructions which operate on the larger, outer-level cache; must first operate on the smaller, inner-level cache. For example, a Hit_Writeback_Invalidate operation targeting the Secondary cache, must first operate on the primary data cache first. If the CACHE instruction implementation does not follow this policy then any software which flushes the caches must mimic this behavior. That is, the software sequences must first operate on the inner cache then operate on the outer cache. The software must place a SYNC instruction after the CACHE instruction whenever there are possible writebacks from the inner cache to ensure that the writeback data is resident in the outer cache before operating on the outer cache. If neither the CACHE instruction implementation nor the software cache flush sequence follow this policy, then the inclusion property of the caches can be broken, which might be a condition that the cache management hardware cannot properly deal with.

When implementing multiple level of caches without the inclusion property, the use of a SYNC instruction after the CACHE instruction is still needed whenever writeback data has to be resident in the next level of memory hierarchy.

For multiprocessor implementations that maintain coherent caches, some of the Hit type of CACHE instruction operations may optionally affect all coherent caches within the implementation. If the effective address uses a coherent

Cache Coherency Attribute (CCA), then the operation is *globalized*, meaning it is broadcast to all of the coherent caches within the system. If the effective address does not use one of the coherent CCAs, there is no broadcast of the operation. If multiple levels of caches are to be affected by one CACHE instruction, all of the affected cache levels must be processed in the same manner - either all affected cache levels use the globalized behavior or all affected cache levels use the non-globalized behavior.

Table 5.19 Encoding of Bits [20:18] of the CACHE Instruction

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b000	I	Index Invalidate	Index	Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire instruction cache by stepping through all valid indices.	Required
	D	Index Writeback Invalidate / Index Invalidate	Index	For a write-back cache: If the state of the cache block at the specified index is valid and dirty, write the block back to the memory address specified by the cache tag. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid.	Required
	S, T	Index Writeback Invalidate / Index Invalidate	Index	For a write-through cache: Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire data cache by stepping through all valid indices. The Index Store Tag must be used to initialize the cache at power up.	Required if S, T cache is implemented
0b001	All	Index Load Tag	Index	Read the tag for the cache block at the specified index into the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. If the <i>DataLo</i> and <i>DataHi</i> registers are implemented, also read the data corresponding to the byte index into the <i>DataLo</i> and <i>DataHi</i> registers. This operation must not cause a Cache Error Exception. The granularity and alignment of the data read into the <i>DataLo</i> and <i>DataHi</i> registers is implementation-dependent, but is typically the result of an aligned access to the cache, ignoring the appropriate low-order bits of the byte index.	Recommended

Table 5.19 Encoding of Bits [20:18] of the CACHE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b010	All	Index Store Tag	Index	Write the tag for the cache block at the specified index from the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. This operation must not cause a Cache Error Exception. This required encoding may be used by software to initialize the entire instruction or data caches by stepping through all valid indices. Doing so requires that the <i>TagLo</i> and <i>TagHi</i> registers associated with the cache be initialized first.	Required
0b011	All	Implementation Dependent	Unspecified	Available for implementation-dependent operation.	Optional
0b100	I, D	Hit Invalidate	Address	If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the instruction cache by stepping through the address range by the line size of the cache.	Required (Instruction Cache Encoding Only), Recommended otherwise
	S, T	Hit Invalidate	Address	In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Optional, if Hit_Invalidate_D is implemented, the S and T variants are recommended.
0b101	I	Fill	Address	Fill the cache from the specified address.	Recommended
	D	Hit Writeback Invalidate / Hit Invalidate	Address	For a write-back cache: If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid.	Required
	S, T	Hit Writeback Invalidate / Hit Invalidate	Address	For a write-through cache: If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the data cache by stepping through the address range by the line size of the cache. In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Required if S, T cache is implemented

Table 5.19 Encoding of Bits [20:18] of the CACHE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b110	D	Hit Writeback	Address	<p>If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After the operation is completed, leave the state of the line valid, but clear the dirty state. For a write-through cache, this operation may be treated as a nop.</p> <p>In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.</p>	Recommended
	S, T	Hit Writeback	Address		Optional, if Hit_Writeback_D is implemented, the S and T variants are recommended.
0b111	I, D	Fetch and Lock	Address	<p>If the cache does not contain the specified address, fill it from memory, performing a write-back if required. Set the state to valid and locked.</p> <p>If the cache already contains the specified address, set the state to locked. In set-associative or fully-associative caches, the way selected on a fill from memory is implementation dependent.</p> <p>The lock state may be cleared by executing an Index Invalidate, Index Writeback Invalidate, Hit Invalidate, or Hit Writeback Invalidate operation to the locked line, or via an Index Store Tag operation to the line that clears the lock bit. Clearing the lock state via Index Store Tag is dependent on the implementation-dependent cache tag and cache line organization, and that Index and Index Writeback Invalidate operations are dependent on cache line organization. Only Hit and Hit Writeback Invalidate operations are generally portable across implementations.</p> <p>It is implementation dependent whether a locked line is displaced as the result of an external invalidate or intervention that hits on the locked line. Software must not depend on the locked line remaining in the cache if an external invalidate or intervention would invalidate the line if it were not locked.</p> <p>It is implementation dependent whether a Fetch and Lock operation affects more than one line. For example, more than one line around the referenced address may be fetched and locked. It is recommended that only the single line containing the referenced address be affected.</p>	Recommended

Restrictions:

The operation of this instruction is **UNDEFINED** for any operation/cache combination that is not implemented. In Release 6, the instruction in this case should perform no operation.

The operation of this instruction is **UNDEFINED** if the operation requires an address, and that address is uncacheable. In Release 6, the instruction in this case should perform no operation.

The operation of the instruction is **UNPREDICTABLE** if the cache line that contains the CACHE instruction is the target of an invalidate or a writeback invalidate.

If this instruction is used to lock all ways of a cache at a specific cache index, the behavior of that cache to subsequent cache misses to that cache index is **UNDEFINED**.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Any use of this instruction that can cause cacheline writebacks should be followed by a subsequent SYNC instruction to avoid hazards where the writeback data is not yet visible at the next level of the memory hierarchy.

This instruction does not produce an exception for a misaligned memory address, since it has no memory access size.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```
vAddr ← GPR[base] + sign_extend(offset)
(pAddr, uncached) ← AddressTranslation(vAddr, DataReadReference)
CacheOp(op, vAddr, pAddr)
```

Exceptions:

TLB Refill Exception.

TLB Invalid Exception

Coprocessor Unusable Exception

Address Error Exception

Cache Error Exception


Bus Error Exception

Programming Notes:

Release 6 architecture implements a 9-bit offset, whereas all release levels lower than Release 6 implement a 16-bit offset.

For cache operations that require an index, it is implementation dependent whether the effective address or the translated physical address is used as the cache index. Therefore, the index value should always be converted to an unmapped address (such as an kseg0 address - by ORing the index with 0x80000000 before being used by the cache instruction). For example, the following code sequence performs a data cache Index Store Tag operation using the index passed in GPR a0:

```
li    a1, 0x80000000    /* Base of kseg0 segment */
or    a0, a0, a1        /* Convert index to kseg0 address */
cache DCIndexStTag, 0(a1) /* Perform the index store tag operation */
```



Format: CACHEE op, offset(base)

microMIPS

Purpose: Perform Cache Operation EVA

To perform the cache operation specified by op using a user mode virtual address while in kernel mode.


Description:

The 9-bit offset is sign-extended and added to the contents of the base register to form an effective address. The effective address is used in one of the following ways based on the operation to be performed and the type of cache as described in the following table.

Table 5.20 Usage of Effective Address

Operation Requires an	Type of Cache	Usage of Effective Address
Address	Virtual	The effective address is used to address the cache. An address translation may or may not be performed on the effective address (with the possibility that a TLB Refill or TLB Invalid exception might occur)
Address	Physical	The effective address is translated by the MMU to a physical address. The physical address is then used to address the cache
Index	N/A	<p>The effective address is translated by the MMU to a physical address. It is implementation dependent whether the effective address or the translated physical address is used to index the cache. As such, a kseg0 address should always be used for cache operations that require an index. See the Programming Notes section below.</p> <p>Assuming that the total cache size in bytes is CS, the associativity is A, and the number of bytes per tag is BPT, the following calculations give the fields of the address which specify the way and the index:</p> $\begin{aligned} \text{OffsetBit} &\leftarrow \text{Log2}(\text{BPT}) \\ \text{IndexBit} &\leftarrow \text{Log2}(\text{CS} / \text{A}) \\ \text{WayBit} &\leftarrow \text{IndexBit} + \text{Ceiling}(\text{Log2}(\text{A})) \\ \text{Way} &\leftarrow \text{Addr}_{\text{WayBit}-1..\text{IndexBit}} \\ \text{Index} &\leftarrow \text{Addr}_{\text{IndexBit}-1..\text{OffsetBit}} \end{aligned}$ <p>For a direct-mapped cache, the Way calculation is ignored and the Index value fully specifies the cache tag. This is shown symbolically in the figure below.</p>

Figure 5.4 Usage of Address Fields to Select Index and Way



A TLB Refill and TLB Invalid (both with cause code equal TLBL) exception can occur on any operation. For index

operations (where the address is used to index the cache but need not match the cache tag) software should use unmapped addresses to avoid TLB exceptions. This instruction never causes TLB Modified exceptions nor TLB Refill exceptions with a cause code of TLBS. This instruction never causes Execute-Inhibit nor Read-Inhibit exceptions.

The effective address may be an arbitrarily-aligned by address. The CACHEE instruction never causes an Address Error Exception due to a non-aligned address.

A Cache Error exception may occur as a by-product of some operations performed by this instruction. For example, if a Writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Similarly, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error. However, cache error exceptions must not be triggered by an Index Load Tag or Index Store tag operation, as these operations are used for initialization and diagnostic purposes.

An Address Error Exception (with cause code equal AdEL) may occur if the effective address references a portion of the kernel address space which would normally result in such an exception. It is implementation dependent whether such an exception does occur.

It is implementation dependent whether a data watch is triggered by a cache instruction whose address matches the Watch register address match conditions.

The CACHEE instruction and the memory transactions which are sourced by the CACHEE instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

Bits [22:21] of the instruction specify the cache on which to perform the operation, as follows:

Table 5.21 Encoding of Bits[22:21] of CACHEE Instruction

Code	Name	Cache
0b00	I	Primary Instruction
0b01	D	Primary Data or Unified Primary
0b10	T	Tertiary
0b11	S	Secondary

Bits [25:23] of the instruction specify the operation to perform. To provide software with a consistent base of cache operations, certain encodings must be supported on all processors. The remaining encodings are recommended

When implementing multiple level of caches and where the hardware maintains the smaller cache as a proper subset of a larger cache, it is recommended that the CACHEE instructions must first operate on the smaller, inner-level cache. For example, a Hit_Writeback_Invalidate operation targeting the Secondary cache, must first operate on the primary data cache first. If the CACHEE instruction implementation does not follow this policy then any software which flushes the caches must mimic this behavior. That is, the software sequences must first operate on the inner cache then operate on the outer cache. The software must place a SYNC instruction after the CACHEE instruction whenever there are possible writebacks from the inner cache to ensure that the writeback data is resident in the outer cache before operating on the outer cache. If neither the CACHEE instruction implementation nor the software cache flush sequence follow this policy, then the inclusion property of the caches can be broken, which might be a condition that the cache management hardware cannot properly deal with.

When implementing multiple level of caches without the inclusion property, you must use SYNC instruction after the CACHEE instruction whenever writeback data has to be resident in the next level of memory hierarchy.

For multiprocessor implementations that maintain coherent caches, some of the Hit type of CACHEE instruction operations may optionally affect all coherent caches within the implementation. If the effective address uses a coherent Cache Coherency Attribute (CCA), then the operation is *globalized*, meaning it is broadcast to all of the coherent

caches within the system. If the effective address does not use one of the coherent CCAs, there is no broadcast of the operation. If multiple levels of caches are to be affected by one CACHEE instruction, all of the affected cache levels must be processed in the same manner — either all affected cache levels use the globalized behavior or all affected cache levels use the non-globalized behavior.

The CACHEE instruction functions the same as the CACHE instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to 1.

Table 5.22 Encoding of Bits [20:18] of the CACHEE Instruction

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b000	I	Index Invalidate	Index	Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire instruction cache by stepping through all valid indices.	Required
	D	Index Writeback Invalidate / Index Invalidate	Index	For a write-back cache: If the state of the cache block at the specified index is valid and dirty, write the block back to the memory address specified by the cache tag. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid.	Required
	S, T	Index Writeback Invalidate / Index Invalidate	Index	For a write-through cache: Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire data cache by stepping through all valid indices. Note that Index Store Tag should be used to initialize the cache at power up.	Required if S, T cache is implemented
0b001	All	Index Load Tag	Index	Read the tag for the cache block at the specified index into the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. If the <i>DataLo</i> and <i>DataHi</i> registers are implemented, also read the data corresponding to the byte index into the <i>DataLo</i> and <i>DataHi</i> registers. This operation must not cause a Cache Error Exception. The granularity and alignment of the data read into the <i>DataLo</i> and <i>DataHi</i> registers is implementation-dependent, but is typically the result of an aligned access to the cache, ignoring the appropriate low-order bits of the byte index.	Recommended

Table 5.22 Encoding of Bits [20:18] of the CACHEE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b010	All	Index Store Tag	Index	Write the tag for the cache block at the specified index from the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. This operation must not cause a Cache Error Exception. This required encoding may be used by software to initialize the entire instruction or data caches by stepping through all valid indices. Doing so requires that the <i>TagLo</i> and <i>TagHi</i> registers associated with the cache be initialized first.	Required
0b011	All	Implementation Dependent	Unspecified	Available for implementation-dependent operation.	Optional
0b100	I, D	Hit Invalidate	Address	If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the instruction cache by stepping through the address range by the line size of the cache.	Required (Instruction Cache Encoding Only), Recommended otherwise
	S, T	Hit Invalidate	Address	In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Optional, if Hit_Invalidate_D is implemented, the S and T variants are recommended.
0b101	I	Fill	Address	Fill the cache from the specified address.	Recommended
	D	Hit Writeback Invalidate / Hit Invalidate	Address	For a write-back cache: If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid.	Required
	S, T	Hit Writeback Invalidate / Hit Invalidate	Address	For a write-through cache: If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the data cache by stepping through the address range by the line size of the cache. In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Required if S, T cache is implemented

Table 5.22 Encoding of Bits [20:18] of the CACHEE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b110	D	Hit Writeback	Address	<p>If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After the operation is completed, leave the state of the line valid, but clear the dirty state. For a write-through cache, this operation may be treated as a nop.</p> <p>In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.</p>	Recommended
	S, T	Hit Writeback	Address		Optional, if Hit_Writeback_D is implemented, the S and T variants are recommended.
0b111	I, D	Fetch and Lock	Address	<p>If the cache does not contain the specified address, fill it from memory, performing a write-back if required. Set the state to valid and locked.</p> <p>If the cache already contains the specified address, set the state to locked. In set-associative or fully-associative caches, the way selected on a fill from memory is implementation dependent.</p> <p>The lock state may be cleared by executing an Index Invalidate, Index Writeback Invalidate, Hit Invalidate, or Hit Writeback Invalidate operation to the locked line, or via an Index Store Tag operation to the line that clears the lock bit. Clearing the lock state via Index Store Tag is dependent on the implementation-dependent cache tag and cache line organization, and that Index and Index Writeback Invalidate operations are dependent on cache line organization. Only Hit and Hit Writeback Invalidate operations are generally portable across implementations.</p> <p>It is implementation dependent whether a locked line is displaced as the result of an external invalidate or intervention that hits on the locked line. Software must not depend on the locked line remaining in the cache if an external invalidate or intervention would invalidate the line if it were not locked.</p> <p>It is implementation dependent whether a Fetch and Lock operation affects more than one line. For example, more than one line around the referenced address may be fetched and locked. It is recommended that only the single line containing the referenced address be affected.</p>	Recommended

Restrictions:

The operation of this instruction is **UNDEFINED** for any operation/cache combination that is not implemented. In Release 6, the instruction in this case should perform no operation.

The operation of this instruction is **UNDEFINED** if the operation requires an address, and that address is uncacheable. In Release 6, the instruction in this case should perform no operation.

The operation of the instruction is **UNPREDICTABLE** if the cache line that contains the CACHEE instruction is the target of an invalidate or a writeback invalidate.

If this instruction is used to lock all ways of a cache at a specific cache index, the behavior of that cache to subsequent cache misses to that cache index is **UNDEFINED**.

Any use of this instruction that can cause cacheline writebacks should be followed by a subsequent SYNC instruction to avoid hazards where the writeback data is not yet visible at the next level of the memory hierarchy.

Only usable when access to Coprocessor0 is enabled and when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

This instruction does not produce an exception for a misaligned memory address, since it has no memory access size.

Operation:

```
vAddr ← GPR[base] + sign_extend(offset)
(pAddr, uncached) ← AddressTranslation(vAddr, DataReadReference)
CacheOp(op, vAddr, pAddr)
```

Exceptions:

TLB Refill Exception.

TLB Invalid Exception

Coprocessor Unusable Exception

Reserved Instruction

Address Error Exception


Cache Error Exception

Bus Error Exception

Programming Notes:

For cache operations that require an index, it is implementation dependent whether the effective address or the translated physical address is used as the cache index. Therefore, the index value should always be converted to a kseg0 address by ORing the index with 0x80000000 before being used by the cache instruction. For example, the following code sequence performs a data cache Index Store Tag operation using the index passed in GPR a0:

```
li    a1, 0x80000000    /* Base of kseg0 segment */
or    a0, a0, a1        /* Convert index to kseg0 address */
cache DCIndexStTag, 0(a1) /* Perform the index store tag operation */
```



Format: CEIL.L.fmt
 CEIL.L.S ft, fs
 CEIL.L.D ft, fs

MIPS64, microMIPS
MIPS64, microMIPS

Purpose: Fixed Point Ceiling Convert to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding up.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and rounding toward $+\infty$ (rounding mode 2). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{\text{NAN2008}}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{\text{NAN2008}}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *fd* for long fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model; it is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


$\text{StoreFPR}(ft, L, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact



Format: CEIL.W.fmt
 CEIL.W.S ft, fs
 CEIL.W.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Ceiling Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding up

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and rounding toward $+\infty$ (rounding mode 2). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{\text{NaN2008}}=0$, the default result is $2^{31}-1$. On cores with $FCSR_{\text{NaN2008}}=1$, the default result is:

- 0 when the input value is NaN
- $2^{31}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{31}-1$
- $-2^{31}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{31}-1$

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:

$\text{StoreFPR}(ft, W, \text{ConvertFmt}(\text{ValueFPR}(fs, \text{fmt}), \text{fmt}, W))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact

31	26	25	21	20	16	15	14	13	6	5	0
POOL32F 010101	rt	fs	00	CFC1 01000000	POOL32FXf 111011						
6	5	5	2	8	6						

Format: CFC1 rt, fs

microMIPS

Purpose: Move Control Word From Floating Point

To copy a word from an FPU control register to a GPR.

Description: $GPR[rt] \leftarrow FP_Control[fs]$

Copy the 32-bit word from FP (coprocessor 1) control register *fs* into GPR *rt*, sign-extending it to 64 bits.

The definition of this instruction has been extended in Release 5 to support user mode read and write of *Status_{FR}* under the control of *Config5_{UFR}*. This optional feature is meant to facilitate transition from *FR*=0 to *FR*=1 floating-point register modes in order to obsolete *FR*=0 mode in a future architecture release. User code may set and clear *Status_{FR}* without kernel intervention, providing kernel explicitly provides permission.

This UFR facility is not supported in Release 6 because Release 6 only allows *FR*=1 mode. Accessing the UFR and UNFR registers causes a Reserved Instruction exception in Release 6 because *FIR_{UFRP}* is always 0.

The definition of this instruction has been extended in Release 6 to allow user code to read and modify the *Config5_{FRE}* bit. Such modification is allowed when this bit is present (as indicated by *FIR_{UFRP}*) and user mode modification of the bit is enabled by the kernel (as indicated by *Config5_{UFE}*). Setting *Config5_{FRE}* to 1 causes all floating point instructions which are not compatible with *FR*=1 mode to take an Reserved Instruction exception. This makes it possible to run pre-Release 6 *FR*=0 floating point code on a Release 6 core which only supports *FR*=1 mode, provided the kernel has been set up to trap and emulate *FR*=0 behavior for these instructions. These instructions include floating-point arithmetic instructions that read/write single-precision registers, LWC1, SWC1, MTC1, and MFC1 instructions.

The FRE facility uses COP1 register aliases FRE and NFRE to access *Config5_{FRE}*.

Restrictions:

There are a few control registers defined for the floating point unit. Prior to Release 6, the result is **UNPREDICTABLE** if *fs* specifies a register that does not exist. In Release 6 and later, a Reserved Instruction exception occurs if *fs* specifies a register that does not exist.

The result is **UNPREDICTABLE** if *fs* specifies the UNFR or NFRE write-only control. Release 6 and later implementations are required to produce a Reserved Instruction exception; software must assume it is **UNPREDICTABLE**.

Operation:

```

if fs = 0 then
    temp ← FIR
elseif fs = 1 then /* read UFR (CP1 Register 1) */
    if FIRUFRP then
        if not Config5UFR then SignalException(ReservedInstruction) endif
        temp ← StatusFR
    else
        if ConfigAR ≥ 2 SignalException(ReservedInstruction) /* Release 6 traps */
        endif
        temp ← UNPREDICTABLE
    endif
endif

```

```

elseif fs = 4 then /* read fs=4 UNFR not supported for reading - UFR suffices */
    if ConfigAR ≥ 2 SignalException(ReservedInstruction) /* Release 6 traps */
    endif
    temp ← UNPREDICTABLE
elseif fs=5 then /* user read of FRE, if permitted */
    if ConfigAR ≤ 2 then temp ← UNPREDICTABLE
    else
        if not Config5UFR then SignalException(ReservedInstruction) endif
        temp ← 031 || Config5FRE
    endif
elseif fs = 25 then /* FCCR */
    temp ← 024 || FCSR31..25 || FCSR23
elseif fs = 26 then /* FEXR */
    temp ← 014 || FCSR17..12 || 05 || FCSR6..2 || 02
elseif fs = 28 then /* FENR */
    temp ← 020 || FCSR11..7 || 04 || FCSR24 || FCSR1..0
elseif fs = 31 then /* FCSR */
    temp ← FCSR
else
    if Config2AR ≥ 2 SignalException(ReservedInstruction)
    /*Release 6 traps; includes NFRE*/
    endif
    temp ← UNPREDICTABLE
endif

if Config2AR < 2 then
    GPR[rt] ← sign_extend(temp)
endif

```

Exceptions:


Coprocessor Unusable, Reserved Instruction

Historical Information:

For the MIPS I, II and III architectures, the contents of GPR *rt* are **UNPREDICTABLE** for the instruction immediately following CFC1.

MIPS V and MIPS32 introduced the three control registers that access portions of FCSR. These registers were not available in MIPS I, II, III, or IV.

MIPS32 Release 5 introduced the UFR and UNFR register aliases that allow user level access to *Status_{FR}*. Release 6 removes them.



Format: CFC2 rt, Impl

microMIPS

The syntax shown above is an example using CFC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Control Word From Coprocessor 2

To copy a word from a Coprocessor 2 control register to a GPR

Description: $GPR[rt] \leftarrow CP2CCR[Impl]$

Copy the 32-bit word from the Coprocessor 2 control register denoted by the *Impl* field, sign-extending it to 64 bits. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The result is **UNPREDICTABLE** if *Impl* specifies a register that does not exist.

Operation:

```
temp ← CP2CCR[Impl]
GPR[rt] ← sign_extend(temp)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: CLASS.fmt
 CLASS fd, fs, fmt
 CLASS.S fd, fs
 CLASS.D fd, fs

microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6

Purpose: Scalar Floating-Point Class Mask

Scalar floating-point class shown as a bit mask for Zero, Negative, Infinite, Subnormal, Quiet NaN, or Signaling NaN.

Description: $FPR[fd] \leftarrow class(FPR[fs])$

Stores in *fd* a bit mask reflecting the floating-point class of the floating point scalar value *fs*.

The mask has 10 bits as follows. Bits 0 and 1 indicate NaN values: signaling NaN (bit 0) and quiet NaN (bit 1). Bits 2, 3, 4, 5 classify negative values: infinity (bit 2), normal (bit 3), subnormal (bit 4), and zero (bit 5). Bits 6, 7, 8, 9 classify positive values: infinity (bit 6), normal (bit 7), subnormal (bit 8), and zero (bit 9).

This instruction corresponds to the **class** operation of the IEEE Standard for Floating-Point Arithmetic 754™-2008. This scalar FPU instruction also corresponds to the vector FCLASS.df instruction of MSA.

The input values and generated bit masks are not affected by the flush-subnormal-to-zero mode FCSR.FS.

The input operand is a scalar value in floating-point data format *fmt*. Bits beyond the width of *fmt* are ignored. The result is a 10-bit bitmask as described above, zero extended to *fmt*-width bits. Coprocessor register bits beyond *fmt*-width bits are UNPREDICTABLE (e.g., for CLASS.S bits 32-63 are UNPREDICTABLE on a 64-bit FPU, while bits 32-128 bits are UNPREDICTABLE if the processor supports MSA).

Restrictions:

No data-dependent exceptions are possible.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

CLASS.fmt is defined only for formats S and D. Other formats must produce a Reserved Instruction exception (unless used for a different instruction).

Operation:

```

if not IsCoprocessorEnabled(1)
    then SignalException(CoprocessorUnusable, 1) endif
if not IsFloatingPointImplemented(fmt)
    then SignalException(ReservedInstruction) endif

fin ← ValueFPR(fs,fmt)
masktmp ← ClassFP(fin, fmt)
StoreFPR (fd, fmt, tmp)
/* end of instruction */

function ClassFP(tt, ts, n)
/* Implementation defined class operation. */
endfunction ClassFP
  
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation

31	26	25	21	20	16	15	6	5	0
POOL32A 000000	rt		rs		CLO 0100101100		POOL32AXf 111100		
6	5		5		10		6		

Format: CLO *rt*, *rs*

microMIPS

Purpose: Count Leading Ones in Word

To count the number of leading ones in a word.

Description: $GPR[rt] \leftarrow \text{count_leading_ones } GPR[rs]$

Bits 31..0 of GPR *rs* are scanned from most significant to least significant bit. The number of leading ones is counted and the result is written to GPR *rt*. If all of bits **31..0** were set in GPR *rs*, the result written to GPR *rt* is 32.

Restrictions:

If GPR *rs* does not contain a sign-extended 32-bit value (bits **63..31** equal), then the results of the operation are **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← 32
for i in 31 .. 0
    if GPR[rs]i = 0 then
        temp ← 31 - i
        break
    endif
endfor
GPR[rt] ← temp

```

Exceptions:

None



Format: CLZ *rt*, *rs*

microMIPS

Purpose: Count Leading Zeros in Word

Count the number of leading zeros in a word.

Description: $GPR[rt] \leftarrow \text{count_leading_zeros } GPR[rs]$

Bits **31..0** of GPR *rs* are scanned from most significant to least significant bit. The number of leading zeros is counted and the result is written to GPR *rt*. If no bits were set in GPR *rs*, the result written to GPR *rt* is 32.

Restrictions:

If GPR *rs* does not contain a sign-extended 32-bit value (bits **63..31** equal), then the results of the operation are **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
temp ← 32
for i in 31 .. 0
    if GPR[rs]i = 1 then
        temp ← 31 - i
        break
    endif
endfor
GPR[rt] ← temp

```

Exceptions:

None

Programming Notes:

Release 6 sets the 'rt' field to a value of 00000.

31	26	25	21	20	16	15	11	10	6	5	4	0													
POOL32F 010101	ft					fs					fd					condn					CMP.condn.S 000101				
POOL32F 010101	ft					fs					fd					condn					CMP.condn.D 010101				
6			5			5			5			5			6										

Format: CMP.condn.fmt
 CMP.condn.S fd, fs, ft
 CMP.condn.D fd, fs, ft

microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Floating Point Compare Setting Mask

To compare FP values and record the result as a format-width mask of all 0s or all 1s in a floating point register

Description: $FPR[fd] \leftarrow FPR[fs] \text{ compare_cond } FPR[ft]$

The value in FPR *fs* is compared to the value in FPR *ft*.

The comparison is exact and neither overflows nor underflows.

If the comparison specified by the *condn* field of the instruction is true for the operand values, the result is true; otherwise, the result is false. If no exception is taken, the result is written into FPR *fd*; true is all 1s and false is all 0s, repeated the operand width of *fmt*. All other bits beyond the operand width *fmt* are UNPREDICTABLE. For example, a 32-bit single precision comparison writes a mask of 32 0s or 1s into bits 0 to 31 of FPR *fd*. It makes bits 32 to 63 UNPREDICTABLE if a 64-bit FPU without MSA is present. It makes bits 32 to 127 UNPREDICTABLE if MSA is present.

The values are in format *fmt*. These instructions, however, do not use an *fmt* field to determine the data type.

The *condn* field of the instruction specifies the nature of the comparison: equals, less than, and so on, unordered or ordered, signalling or quiet, as specified in [Table 5.23 “Comparing CMP.condn.fmt, IEEE 754-2008, C.cond.fmt, and MSA FP compares” on page 181](#).

Release 6: The *condn* field bits have specific purposes: *cond₄* and *cond_{2..1}* specify the nature of the comparison (equals, less than, and so on); *cond₀* specifies whether the comparison is ordered or unordered, that is false or true if any operand is a NaN; *cond₃* indicates whether the instruction should signal an exception on QNaN inputs. However, in the future the MIPS ISA may be extended in ways that do not preserve these meanings.

All encodings of the *condn* field that are not specified (for example, items shaded in [Table 5.23](#)) are reserved in Release 6 and produce a Reserved Instruction exception.

If one of the values is an SNaN, or if a signalling comparison is specified and at least one of the values is a QNaN, an Invalid Operation condition is raised and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written and an Invalid Operation exception is taken immediately. Otherwise, the mask result is written into FPR *fd*.

There are four mutually exclusive ordering relations for comparing floating point values; one relation is always true and the others are false. The familiar relations are *greater than*, *less than*, and *equal*. In addition, the IEEE floating point standard defines the relation *unordered*, which is true when at least one operand value is NaN; NaN compares unordered with everything, including itself. Comparisons ignore the sign of zero, so +0 equals -0.

The comparison condition is a logical predicate, or equation, of the ordering relations such as *less than or equal*, *equal*, *not less than*, or *unordered or equal*. Compare distinguishes among the 16 comparison predicates. The Boolean result of the instruction is obtained by substituting the Boolean value of each ordering relation for the two FP values in the equation. For example: If the *equal* relation is true, then all four example predicates above yield a true result. If the *unordered* relation is true then only the final predicate, *unordered or equal*, yields a true result.

The predicates implemented are described in [Table 5.23 “Comparing CMP.condn.fmt, IEEE 754-2008, C.cond.fmt, and MSA FP compares” on page 181](#). Not all of the 16 IEEE predicates are implemented directly by hardware. For the directed comparisons (LT, LE, GT, GE) the missing predicates can be obtained by reversing the FPR register operands *ft* and *fs*. For example, the hardware implements the “Ordered Less Than” predicate LT(*fs*,*ft*); reversing the operands LT(*ft*,*fs*) produces the dual predicate “Unordered or Greater Than or Equal” UGE(*fs*,*ft*). [Table 5.23](#) shows these mappings. Reversing inputs is ineffective for the symmetric predicates such as EQ; Release 6 implements these negative predicates directly, so that all mask values can be generated in a single instruction.

[Table 5.23](#) compares CMP.condn.fmt to (1) the MIPS32 Pre-Release 6 C.cond.fmt instructions, and (2) the (MSA) MIPS SIMD Architecture packed vector floating point comparison instructions. CMP.condn.fmt provides exactly the same comparisons for FPU scalar values that MSA provides for packed vectors, with similar mnemonics. CMP.condn.fmt provides a superset of the MIPS32 Release 5 C.cond.fmt comparisons.

In addition, [Table 5.23](#) shows the corresponding IEEE 754-2008 comparison operations.

Table 5.23 Comparing CMP.condn.fmt, IEEE 754-2008, C.cond.fmt, and MSA FP compares

Shaded entries in the table are unimplemented, and reserved.

Instruction Encodings																								
CMP.condn.fmt: 010001 fffff ttttt sssss dddd 0cccc C.cond.fmt: 010001 fffff ttttt sssss CCC00 11cccc MSA: 011110 000of ttttt sssss dddd mmmmm																								
Invalid Operand Exception	MSA: operation oooo Bits 25...22 C: cond cccc - Bits 3..0 CMP: condn cccccc - Bits 3..0		MSA: minor opcode mmmmm Bits 5...0 = 26 - 011010 CMP: condn Bit 5..4 = 00 C: only applicable										MSA: minor opcode mmmmm Bits 5...0 = 28 - 011100 CMP: condn Bit 5..4 = 01 C: not applicable											
			Predicates								Negated Predicates													
			Relation > < = ? C condn.fmt				MSA	CMP condn.fmt	Long names		IEEE		Relation > < = ? C condn.fmt				MSA	CMP condn.fmt	Long names		IEEE			
no (non-signalling) yes (always signal NaN)	0	0000	F	F	F	F	F	FCAF	AF	False Always False				T	T	T	T	T		AT	True Always True			
	1	0001	F	F	F	T	UN	FCUN	UN	Unordered		compareQuietUnordered ? isUnordered		T	T	T	F	OR	FCOR	OR	Ordered	compareQuietOrdered <=> NOT(isUnordered)		
	2	0010	F	F	T	F	EQ	FCEQ	EQ	Equal		compareQuietEqual =		T	T	F	T	NEQ	FCUNE	UNE	Not Equal	compareQuietNotEqual ?<>, NOT(=), ≠		
	3	0011	F	F	T	T	UEQ	FCUEQ	UEQ	Unordered or Equal				T	T	F	F	OGL	FCNE	NE	Ordered Greater Than or Less Than			
	4	0100	F	T	F	F	OLT	FCLT	LT	Ordered Less Than		compareQuietLess isLess		T	F	T	T	UGE		UGE	Unordered or Greater Than or Equal	compareQuietNotLess ?>=, NOT(isLess)		
	5	0101	F	T	F	T	ULT	FCULT	ULT	Unordered or Less Than		compareQuietLessUnor- dered ?<, NOT(isGreaterEqual)		T	F	T	F	OGE		OGE	Ordered Greater Than or Equal	compareQuiet- GreatrEqual isGreaterEqual		
	6	0110	F	T	T	F	OLE	FCLE	LE	Ordered Less than or Equal		compareQuietLessEqual isLessEqual		T	F	F	T	UGT		UGT	Unordered or Greater Than	compareQuietGreaterUn- ordered ?>, NOT(isLessEqual)		
	7	0111	F	T	T	T	ULE	FCULE	ULE	Unordered or Less Than or Equal		compareQuietNotGreater ?<=, NOT(isGreater)		T	F	F	F	OGT		OGT	Ordered Greater Than	compareQuietGreater isGreater		

Table 5.23 Comparing CMP.condn.fmt, IEEE 754-2008, C.cond.fmt, and MSA FP compares (Continued)

Shaded entries in the table are unimplemented, and reserved.

Instruction Encodings																					
CMP.condn.fmt: 010001 fffff ttttt sssss dddd 0cccc C.cond.fmt: 010001 fffff ttttt sssss CCC00 11cccc MSA: 011110 000of ttttt sssss dddd mmmmm																					
Invalid Operand Exception		MSA: minor opcode mmmmm Bits 5...0 = 26 - 011010 CMP: condn Bit 5.4 = 00 C: only applicable										MSA: minor opcode mmmmm Bits 5...0 = 28 - 011100 CMP: condn Bit 5.4 = 01 C: not applicable									
		Predicates										Negated Predicates									
		Relation				C condn.fmt	MSA	CMP condn.fmt	Long names	IEEE	Relation				C condn.fmt	MSA	CMP condn.fmt	Long names	IEEE		
>	<	=	?	>	<						=	?									
yes (signalling)		8	1000	F	F	F	F	SF	FSAF	SAF	Signalling False Signalling Always False		T	T	T	T	ST		SAT	Signalling True Signalling Always True	
		9	1001	F	F	F	T	NGLE	FSUN	SUN	Not Greater Than or Less Than or Equal Signalling Unordered		T	T	T	F	GLE	FSOR	SOR	Greater Than or Less Than or Equal Signalling Ordered	
		10	1010	F	F	T	F	SEQ	FSEQ	SEQ	Signalling Equal Ordered Signalling Equal	compareSignalling Equal	T	T	F	T	SNE	FSUNE	SUNE	Signalling Not Equal Signalling Unor- dered or Not Equal	compareSignalling- NotEqual
		11	1011	F	F	T	T	NGL	FSUEQ	SUEQ	Not Greater Than or Less Than Signalling Unordered or Equal		T	T	F	F	GL	FSNE	SNE	Greater Than or Less Than Signalling Ordered Not Equal	
		12	1100	F	T	F	F	LT	FSLT	SLT	Less Than Ordered Signalling Less Than	compareSignallingLess <	T	F	T	T	NLT		SUGE	Not Less Than Signalling Unordered or Greater Than or Equal	compareSignallingNot- Less NOT(<=)
		13	1101	F	T	F	T	NGE	FSULT	SULT	Not Greater Than or Equal Unordered or Less Than	compareSignalling- LessUnordered NOT(>=)	T	F	T	F	GE		SUGE	Signalling Ordered Greater Than or Equal	compareSignalling- GreaterEqual >=, ≥
		14	1110	F	T	T	F	LE	FSLE	SLE	Less Than or Equal Ordered Signalling Less Than or Equal	compareSignalling- LessEqual <=, ≤	T	F	F	T	NLE		SUGT	Not Less Than or Equal Signalling Unordered or Greater Than	compareSignalling- GreaterUnordered NOT(<=)
		15	1111	F	T	T	T	NGT	FSULE	SULE	Not Greater Than Signalling Unordered or Less Than or Equal	compareSignalling- NotGreater NOT(>)	T	F	F	F	GT		SOGT	Greater Than Signalling Ordered Greater Than	compareSignalling- Greater >

Restrictions:**Operation:**

```

if SNaN(ValueFPR(fs, fmt)) or SNaN(ValueFPR(ft, fmt)) or
   QNaN(ValueFPR(fs, fmt)) or QNaN(ValueFPR(ft, fmt))
then
  less ← false
  equal ← false
  unordered ← true
  if (SNaN(ValueFPR(fs,fmt)) or SNaN(ValueFPR(ft,fmt))) or
     (cond3 and (QNaN(ValueFPR(fs,fmt)) or QNaN(ValueFPR(ft,fmt)))) then
    SignalException(InvalidOperation)
  endif
else
  less ← ValueFPR(fs, fmt) <fmt ValueFPR(ft, fmt)
  equal ← ValueFPR(fs, fmt) =fmt ValueFPR(ft, fmt)
  unordered ← false
endif
condition ← cond4 xor (
  (cond2 and less)
  or (cond1 and equal)
  or (cond0 and unordered) )

StoreFPR (fd, fmt, ExtendBit.fmt(condition))

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation

31	26	25		3	2	0
POOL32A 000000			cofun			COP2 010
6			23			3

Format: COP2 func

microMIPS

Purpose: Coprocessor Operation to Coprocessor 2

To perform an operation to Coprocessor 2.

Description: `CoprocessorOperation(2, cofun)`

An implementation-dependent operation is performed to Coprocessor 2, with the *cofun* value passed as an argument. The operation may specify and reference internal coprocessor registers, and may change the state of the coprocessor conditions, but does not modify state within the processor. Details of coprocessor operation and internal state are described in the documentation for each Coprocessor 2 implementation.


Restrictions:

Operation:

`CoprocessorOperation(2, cofun)`

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: CRC32B, CRC32H, CRC32W, CRC32D
 CRC32B rt, rs, rt
 CRC32H rt, rs, rt
 CRC32W rt, rs, rt
 CRC32D rt, rs, rt

microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS64 Release 6

Purpose: Generate CRC with reversed polynomial 0xEDB88320

Description: $\text{GPR}[\text{rt}] \leftarrow \text{CRC32}(\text{GRP}[\text{rs}], \text{GPR}[\text{rt}])$

CRC32B/H/W/D generates a 32-bit Cyclic Redundancy Check (CRC) value based on the reversed polynomial 0xEDB88320. The new 32-bit CRC value is generated with a cumulative 32-bit CRC value input as GPR[rt] and a byte or half-word or word or double-word message right-justified in GPR[rs]. The message size is encoded in field sz of the instruction.

The generated value overwrites the input CRC value in GPR[rt], after sign extension, as the original value is considered redundant once the cumulative CRC value is re-generated with the additional message. More importantly, source-destroying definition of the CRC instruction allows the instruction to be included in a loop without having to move the destination to the source for the next iteration of the instruction.

The CRC32B/H/W/D instruction does not pad the input message. It is software's responsibility to ensure the input message, whether byte, half-word, word or double-word is fully-defined, otherwise the result is UNPREDICTABLE and thus unusable.

The reversed polynomial is a 33-bit polynomial of degree 32. Since the coefficient of most significance is always 1, it is dropped from the 32-bit binary number, as per standard representation. The order of the remaining coefficients increases from right to left in the binary representation.

Since the CRC is processed more than a bit at a time, the order of bits in the data elements of size byte, half-word, word or double-word is important. The specification assumes support for an "lsb-first" (little-endian) standard, and thus coefficients of polynomial terms that represent the message must be ordered from right to left in order of decreasing significance.

The specification of the CRC instruction assumes the following in regards to a message of arbitrary length whose 32-bit CRC value is to be generated. The message itself is a polynomial represented by binary coefficients of each term of the polynomial.

- The message is a sequence of bytes or half-words or words or double-words as per use case. The appropriate instruction is chosen.
- For each message element of size byte/half-word/word/double-word, the least-significant bit corresponds to the most significant coefficient, and significance decreases from right to left.
- Message elements themselves must be processed in order of decreasing significance, with reference to coefficients of the terms of the polynomial the message represents.
- The polynomial is thus reversed to match the order of coefficients for the message of arbitrary length.
- The resultant CRC is also a polynomial whose coefficients are arranged in decreasing significance from right to left.

The typical use of CRC is to generate a checksum to accompany a message that is transmitted electronically in order to detect transmission errors at the receiving end. If the message is considered to be a polynomial, the coefficient of the most-significant term is transmitted first followed by remaining bits in order of decreasing significance, followed by the 32-bit CRC. The specification for these CRC instruction is thus most appropriate for standards that transmit least significant bit of data first (little-endian), such as IEEE 802 Ethernet. The least-significant bit of data conveniently maps to the coefficient of the most-significant term of the message polynomial.

Restrictions:

No data-dependent exceptions are possible.

Operation:

```

if (Config5CRCP = 0) then
    SignalException(ReservedInstruction)
endif

if (sz = 0b00) then
    temp ← CRC32(GPR[rt], GPR[rs], 1, 0xEDB88320)
else if (sz = 0b01) then
    temp ← CRC32(GPR[rt], GPR[rs], 2, 0xEDB88320)
else if (sz = 0b10) then
    temp ← CRC32(GPR[rt], GPR[rs], 4, 0xEDB88320)
else if (sz = 0b11) then
    if (Are64BitOperationsEnabled()) then
        temp ← CRC32(GPR[rt], GPR[rs], 8, 0xEDB88320)
    else
        SignalException(ReservedInstruction)
    endif
endif
endif
GPR[rt] ← sign_extend(temp)

// Bit oriented definition of CRC32 function
function CRC32(value, message, numbytes, poly)
    # value - right-justified current 32-bit CRC value
    # message - right-justified byte/half-word/word/double-word message
    # numbytes - size of message in bytes: byte/half-word/word/double-word
    # poly - 32-bit reversed polynomial

    value ← (value and 0xffffffff) xor {(64-(numbytes*8))'b0,message}
    for (i=0; i<numbytes*8; i++)
        if (value and 0d1) then // check most significant coefficient
            value ← (value >> 1) xor poly
        else
            value ← (value >> 1)
        endif
    endfor
    return value
endfunction

```

Exceptions:

Reserved Instruction Exception

Restriction:

These instructions are implemented in Release 6 only if CP0 *Config5_{CRCP}* is set to 1.

Programing Notes:

When calculating CRC, it is recommended that the initial value of GPR[rt] be all ones, when the CRC instruction is the first in the sequence to be referenced. This allows the CRC to differentiate between actual leading zeroes in the message element, and zeros added by transmission errors. The initial all one's value makes no difference to the CRC calculation as long as both sender and receiver use the same assumption on the initial value, to generate and check respectively.

If the order of bits in bytes assumes most-significant bit first, then Release 6 BITSWAP can be used to reverse the order of bits in order to operate with these instructions. However BITSWAP would only apply to byte messages.

CRC32B/H/W/D instructions are interchangeable: a series of low-order CRC instructions can be reduced to a series of high-order CRC32H operations, to increase throughput of the overall CRC generation process. The process of doing this will add trailing zeroes to the message for which CRC is being generated, since the data element is now larger, however, this will not change the CRC value that is generated. It is the original message that must be transmitted along with the CRC, without the trailing zeroes.

In pseudo-assembly, the following sequence of byte CRC operations may be used to generate a cumulative CRC value. (Pseudo-assembly is used to clearly indicate terms which need to be modified for interchangeability.)


```
li $3, 0xFFFF_FFFF    // initialize CRC value
la $4, memaddr         // assume word-aligned for convenience

for (i=0; i < byte_cnt; i++)
    lb $2, 0($4)        // read message bytes
    crc32b $3, $2, $3
    add $4, $4, 1        // increment byte memory address by 1
endfor
```

This is equivalent to the sequence of word CRC operations. The simple example assumes some multiple of 4 bytes are processed.

```
for (i=0; i < byte_cnt/4; i++)
    lw $2, 0($4)         // read message words
    crc32w $3, $2, $3
    add $4, $4, 4        // increment word memory address by 4
endfor
```

The throughput is thus increased by a multiple of 4 as only a quarter of the byte oriented operations occur.



Format: CRC32CB, CRC32CH, CRC32CW, CRC32CD

CRC32CB rt, rs, rt

CRC32CH rt, rs, rt

CRC32CW rt, rs, rt

CRC32CD rt, rs, rt

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS32 Release 6

microMIPS64 Release 6

Purpose: Generate CRC with reversed polynomial 0x82F63B78

Description: $\text{GPR}[\text{rt}] \leftarrow \text{CRC32C}(\text{GPR}[\text{rs}], \text{GPR}[\text{rt}])$

CRC32CB/H/W/D generates a 32-bit Cyclic Redundancy Check (CRC) value based on the reversed polynomial 0x82F63B78 (Castagnoli). The new 32-bit CRC value is generated with a cumulative 32-bit CRC value input as GPR[rt] and a byte or half-word or word or double-word message right-justified in GPR[rs]. The message size is encoded in field sz of the instruction.

The generated value overwrites the input CRC value in GPR[rt], after sign extension, as the original value is considered redundant once the cumulative CRC value is re-generated with the additional message. More importantly, source-destroying definition of the CRC instruction allows the instruction to be included in a loop without having to move the destination to the source for the next iteration of the instruction.

The CRC32CB/H/W/D instruction does not pad the input message. It is software's responsibility to ensure the input message, whether byte, half-word, word or double-word is fully-defined, otherwise the result is UNPREDICTABLE and thus unusable.

The reversed polynomial is a 33-bit polynomial of degree 32. Since the coefficient of most significance is always 1, it is dropped from the 32-bit binary number, as per standard representation. The order of the remaining coefficients increases from right to left in the binary representation.

Since the CRC is processed more than a bit at a time, the order of bits in the data elements of size byte, half-word, word or double-word is important. The specification assumes support for an "lsb-first" (little-endian) standard, and thus coefficients of polynomial terms that represent the message must be ordered from right to left in order of decreasing significance.

The specification of the CRC instruction assumes the following in regards to a message of arbitrary length whose 32-bit CRC value is to be generated. The message itself is a polynomial represented by binary coefficients of each term of the polynomial.

- The message is a sequence of bytes or half-words or words or double-words as per use case. The appropriate instruction is chosen.
- For each message element of size byte/half-word/word/double-word, the least-significant bit corresponds to the most significant coefficient, and significance decreases from right to left.
- Message elements themselves must be processed in order of decreasing significance, with reference to coefficients of the terms of the polynomial the message represents.
- The polynomial is thus reversed to match the order of coefficients for the message of arbitrary length.
- The resultant CRC is also a polynomial whose coefficients are arranged in decreasing significance from right to left.

The typical use of CRC is to generate a checksum to accompany a message that is transmitted electronically in order to detect transmission errors at the receiving end. If the message is considered to be a polynomial, the coefficient of the most-significant term is transmitted first followed by remaining bits in order of decreasing significance, followed by the 32-bit CRC. The specification for these CRC instruction is thus most appropriate for standards that transmit least significant bit of data first (little-endian), such as IEEE 802 Ethernet. The least-significant bit of data conveniently maps to the coefficient of the most-significant term of the message polynomial.

Restrictions:

No data-dependent exceptions are possible.

Operation:

```

if (Config5CRCP = 0) then
    SignalException(ReservedInstruction)
endif

if (sz = 0b00) then
    temp ← CRC32(GPR[rt], GPR[rs], 1, 0x82F63B78)
else if (sz = 0b01) then
    temp ← CRC32(GPR[rt], GPR[rs], 2, 0x82F63B78)
else if (sz = 0b10) then
    temp ← CRC32(GPR[rt], GPR[rs], 4, 0x82F63B78)
else if (sz = 0b11) then
    if (Are64BitOperationsEnabled()) then
        temp ← CRC32(GPR[rt], GPR[rs], 8, 0x82F63B78)
    else
        SignalException(ReservedInstruction)
    endif
endif
endif
GPR[rt] ← sign_extend(temp)

// Bit oriented definition of CRC32 function
function CRC32(value, message, numbytes, poly)
    # value - right-justified current 32-bit CRC value
    # message - right-justified byte/half-word/word/double-word message
    # numbytes - size of message in bytes: byte/half-word/word/double-word
    # poly - 32-bit reversed polynomial

    value ← (value and 0xffffffff) xor {(64-(numbytes*8))'b0,message}
    for (i=0; i<numbytes*8; i++)
        if (value and 0d1) then // check most significant coefficient
            value ← (value >> 1) xor poly
        else
            value ← (value >> 1)
        endif
    endfor
    return value
endfunction

```

Exceptions:

Reserved Instruction Exception

Restriction:

These instructions are implemented in Release 6 only if CP0 *Config5_{CRCP}* is set to 1.

Programing Notes:

When calculating CRC, it is recommended that the initial value of GPR[rt] be all ones, when the CRC instruction is the first in the sequence to be referenced. This allows the CRC to differentiate between actual leading zeroes in the message element, and zeros added by transmission errors. The initial all one's value makes no difference to the CRC calculation as long as both sender and receiver use the same assumption on the initial value, to generate and check respectively.

If the order of bits in bytes assumes most-significant bit first, then Release 6 BITSWAP can be used to reverse the order of bits in order to operate with these instructions. However BITSWAP would only apply to byte messages.

CRC32CB/H/W/D instructions are interchangeable: a series of low-order CRC instructions can be reduced to a series of high-order CRC32CH operations, to increase throughput of the overall CRC generation process. The process of doing this will add trailing zeroes to the message for which CRC is being generated, since the data element is now larger, however, this will not change the CRC value that is generated. It is the original message that must be transmitted along with the CRC, without the trailing zeroes.

In pseudo-assembly, the following sequence of byte CRC operations may be used to generate a cumulative CRC value. (Pseudo-assembly is used to clearly indicate terms which need to be modified for interchangeability.)


```
li $3, 0xFFFF_FFFF    // initialize CRC value
la $4, memaddr         // assume word-aligned for convenience

for (i=0; i < byte_cnt; i++)
    lb $2, 0($4)        // read message bytes
    crc32cb $3, $2, $3
    add $4, $4, 1        // increment byte memory address by 1
endfor
```

This is equivalent to the sequence of word CRC operations. The simple example assumes some multiple of 4 bytes are processed.

```
for (i=0; i < byte_cnt/4; i++)
    lw $2, 0($4)        // read message words
    crc32cw $3, $2, $3
    add $4, $4, 4        // increment word memory address by 4
endfor
```

The throughput is thus increased by a multiple of 4 as only a quarter of the byte oriented operations occur.



Format: CTC1 rt, fs

microMIPS

Purpose: Move Control Word to Floating Point

To copy a word from a GPR to an FPU control register.

Description: $FP_Control[fs] \leftarrow GPR[rt]$

Copy the low word from GPR *rt* into the FP (coprocessor 1) control register indicated by *fs*.

Writing to the floating point *Control/Status* register, the *FCSR*, causes the appropriate exception if any *Cause* bit and its corresponding *Enable* bit are both set. The register is written before the exception occurs. Writing to *FEXR* to set a cause bit whose enable bit is already set, or writing to *FENR* to set an enable bit whose cause bit is already set causes the appropriate exception. The register is written before the exception occurs and the *EPC* register contains the address of the CTC1 instruction.

The definition of this instruction has been extended in Release 5 to support user mode read and write of *Status_{FR}* under the control of *Config5_{UFR}*. This optional feature is meant to facilitate transition from *FR*=0 to *FR*=1 floating-point register modes in order to obsolete *FR*=0 mode in a future architecture release. User code may set and clear *Status_{FR}* without kernel intervention, providing kernel explicitly provides permission.

This UFR facility is not supported in Release 6 since Release 6 only allows *FR*=1 mode. Accessing the UFR and UNFR registers causes a Reserved Instruction exception in Release 6 since *FIR_{UFRP}* is always 0.

The definition of this instruction has been extended in Release 6 to allow user code to read and modify the *Config5_{FRE}* bit. Such modification is allowed when this bit is present (as indicated by *FIR_{UFRP}*) and user mode modification of the bit is enabled by the kernel (as indicated by *Config5_{UFE}*). Setting *Config5_{FRE}* to 1 causes all floating point instructions which are not compatible with *FR*=1 mode to take an Reserved Instruction exception. This makes it possible to run pre-Release 6 *FR*=0 floating point code on a Release 6 core which only supports *FR*=1 mode, provided the kernel has been set up to trap and emulate *FR*=0 behavior for these instructions. These instructions include floating-point arithmetic instructions that read/write single-precision registers, LWC1, SWC1, MTC1, and MFC1 instructions.

The FRE facility uses COP1 register aliases FRE and NFRE to access *Config5_{FRE}*.

Restrictions:

There are a few control registers defined for the floating point unit. Prior to Release 6, the result is **UNPREDICTABLE** if *fs* specifies a register that does not exist. In Release 6 and later, a Reserved Instruction exception occurs if *fs* specifies a register that does not exist.

Furthermore, the result is **UNPREDICTABLE** if *fd* specifies the UFR, UNFR, FRE and NFRE aliases, with *fs* anything other than 00000, GPR[0]. Release 6 implementations and later are required to produce a Reserved Instruction exception; software must assume it is **UNPREDICTABLE**.

Operation:

```
temp ← GPR[rt]31..0
if (fs = 1 or fs = 4) then
    /* clear UFR or UNFR(CP1 Register 1) */
    if ConfigAR ≥ 2 SignalException(ReservedInstruction) /* Release 6 traps */ endif
```



```

    if not Config5UFR then SignalException(ReservedInstruction) endif
    if not (rt = 0 and FIRUFRP) then UNPREDICTABLE /*end of instruction*/ endif
    if fs = 1 then StatusFR ← 0
    elseif fs = 4 then StatusFR ← 1
    else /* cannot happen */
elseif fs=5 then /* user write of 1 to FRE, if permitted */
    if ConfigAR ≤ 2 then UNPREDICTABLE
    else
        if rt ≠ 0 then SignalException(ReservedInstruction) endif
        if not Config5UFR then SignalException(ReservedInstruction) endif
        Config5UFR ← 0
    endif
elseif fs=6 then /* user write of 0 to FRE, if permitted (NFRE alias) */
    if ConfigAR ≤ 2 then UNPREDICTABLE
    else
        if rt ≠ 0 then SignalException(ReservedInstruction) endif
        if not Config5UFR then SignalException(ReservedInstruction) endif
        Config5UFR ← 1
    endif
elseif fs = 25 then /* FCCR */
    if temp31..8 ≠ 024 then
        UNPREDICTABLE
    else
        FCSR ← temp7..1 || FCSR24 || temp0 || FCSR22..0
    endif
elseif fs = 26 then /* FEXR */
    if temp31..18 ≠ 0 or temp11..7 ≠ 0 or temp2..0 ≠ 0 then
        UNPREDICTABLE
    else
        FCSR ← FCSR31..18 || temp17..12 || FCSR11..7 ||
            temp6..2 || FCSR1..0
    endif
elseif fs = 28 then /* FENR */
    if temp31..12 ≠ 0 or temp6..3 ≠ 0 then
        UNPREDICTABLE
    else
        FCSR ← FCSR31..25 || temp2 || FCSR23..12 || temp11..7
            || FCSR6..2 || temp1..0
    endif
elseif fs = 31 then /* FCSR */
    if (FCSRImpl field is not implemented) and(temp22..18 ≠ 0) then
        UNPREDICTABLE
    elseif (FCSRImpl field is implemented) and temp20..18 ≠ 0 then
        UNPREDICTABLE
    else
        FCSR ← temp
    endif
else
    if Config2AR ≥ 2 SignalException(ReservedInstruction) /* Release 6 traps */
    endif
    UNPREDICTABLE
endif
CheckFPException()

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Division-by-zero, Inexact, Overflow, Underflow


Historical Information:

For the MIPS I, II and III architectures, the contents of floating point control register *fs* are **UNPREDICTABLE** for the instruction immediately following CTC1.

MIPS V and MIPS32 introduced the three control registers that access portions of *FCSR*. These registers were not available in MIPS I, II, III, or IV.

MIPS32 Release 5 introduced the UFR and UNFR register aliases that allow user level access to *Status_{FR}*.

MIPS32 Release 6 introduced the FRE and NFRE register aliases that allow user to cause traps for *FR*=0 mode emulation.



Format: CTC2 *rt*, Impl

microMIPS

The syntax shown above is an example using CTC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Control Word to Coprocessor 2

To copy a word from a GPR to a Coprocessor 2 control register.

Description: $CP2CCR[Impl] \leftarrow GPR[rt]$

Copy the low word from GPR *rt* into the Coprocessor 2 control register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The result is **UNPREDICTABLE** if *rd* specifies a register that does not exist.

Operation:

```
temp ← GPR[rt]31..0
CP2CCR[Impl] ← temp
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: CVT.D.fmt
 CVT.D.S ft, fs
 CVT.D.W ft, fs
 CVT.D.L ft, fs

microMIPS
 microMIPS
 MIPS64, microMIPS

Purpose: Floating Point Convert to Double Floating Point

To convert an FP or fixed point value to double FP.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in double floating point format and rounded according to the current rounding mode in *FCSR*. The result is placed in FPR *ft*. If *fmt* is S or W, then the operation is always exact.

Restrictions:

The fields *fs* and *ft* must specify valid FPRs, *fs* for type *fmt* and *ft* for double floating point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

For CVT.D.L, the result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model.

Operation:


`StoreFPR (ft, D, ConvertFmt(ValueFPR(fs, fmt), fmt, D))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact



Format: CVT.L.fmt
 CVT.L.S ft, fs
 CVT.L.D ft, fs

MIPS64, microMIPS
MIPS64, microMIPS

Purpose: Floating Point Convert to Long Fixed Point

To convert an FP value to a 64-bit fixed point.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

Convert the value in format *fmt* in FPR *fs* to long fixed point format and round according to the current rounding mode in *FCSR*. The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range $-2^{63}-1$ to $2^{63}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs, *fs* for type *fmt* and *fd* for long fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model; it is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


StoreFPR (ft, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact,



Format: CVT.S.fmt
 CVT.S.D ft, fs
 CVT.S.W ft, fs
 CVT.S.L ft, fs

microMIPS
 microMIPS
 MIPS64, microMIPS

Purpose: Floating Point Convert to Single Floating Point

To convert an FP or fixed point value to single FP.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in single floating point format and rounded according to the current rounding mode in *FCSR*. The result is placed in FPR *ft*.

Restrictions:

The fields *fs* and *ft* must specify valid FPRs—*fs* for type *fmt* and *fd* for single floating point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

For CVT.S.L, the result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model; it is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


$\text{StoreFPR}(ft, S, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, S))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow, Underflow



Format: CVT.W.fmt
 CVT.W.S ft, fs
 CVT.W.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and rounded according to the current rounding mode in *FCSR*. The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *ft* for word fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


$\text{StoreFPR}(ft, W, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, W))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact



Format: DADD rd, rs, rt

microMIPS64

Purpose: Doubleword Add

To add 64-bit integers. If overflow occurs, then trap.

Description: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

The 64-bit doubleword value in GPR *rt* is added to the 64-bit value in GPR *rs* to produce a 64-bit result. If the addition results in 64-bit 2's complement arithmetic overflow, then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 64-bit result is placed into GPR *rd*.

Restrictions:

Operation:


```
temp ← (GPR[rs]63 || GPR[rs]) + (GPR[rt]63 || GPR[rt])
if (temp64 ≠ temp63) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp63..0
endif
```

Exceptions:

Integer Overflow, Reserved Instruction

Programming Notes:

DADDU performs the same arithmetic operation but does not trap on overflow.



Format: DADDIU *rt*, *rs*, *immediate*

microMIPS64

Purpose: Doubleword Add Immediate Unsigned

To add a constant to a 64-bit integer

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] + \text{sign_extend}(\text{immediate})$

The 16-bit signed *immediate* is added to the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rt*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

Operation:


$\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] + \text{sign_extend}(\text{immediate})$

Exceptions:

Reserved Instruction

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for unsigned arithmetic such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: DADDU rd, rs, rt

microMIPS64

Purpose: Doubleword Add Unsigned

To add 64-bit integers

Description: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

The 64-bit doubleword value in GPR *rt* is added to the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

Operation:


$$GPR[rd] \leftarrow GPR[rs] + GPR[rt]$$

Exceptions:

Reserved Instruction

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for unsigned arithmetic such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: DCLO *rd*, *rs*

microMIPS64

Purpose: Count Leading Ones in Doubleword

To count the number of leading ones in a doubleword

Description: $\text{GPR}[\text{rd}] \leftarrow \text{count_leading_ones } \text{GPR}[\text{rs}]$

The 64-bit word in GPR *rs* is scanned from most-significant to least-significant bit. The number of leading ones is counted and the result is written to GPR *rd*. If all 64 bits were set in GPR *rs*, the result written to GPR *rd* is 64.

Restrictions:


To be compliant with the MIPS32 and MIPS64 Architecture, software must place the same GPR number in both the *rt* and *rd* fields of the instruction. The operation of the instruction is **UNPREDICTABLE** if the *rt* and *rd* fields of the instruction contain different values. Release 6's new instruction encoding does not contain an field; Release 6 implementations are required to signal a Reserved Instruction exception if the field is nonzero.

Operation:

```
temp ← 64
for i in 63.. 0
  if GPR[rs]i = 0 then
    temp ← 63 - i
    break
  endif
endfor
GPR[rd] ← temp
```

Exceptions:

None



Format: DCLZ *rd*, *rs*

microMIPS64

Purpose: Count Leading Zeros in Doubleword

To count the number of leading zeros in a doubleword

Description: $\text{GPR}[\text{rd}] \leftarrow \text{count_leading_zeros } \text{GPR}[\text{rs}]$

The 64-bit word in GPR *rs* is scanned from most significant to least significant bit. The number of leading zeros is counted and the result is written to GPR *rd*. If no bits were set in GPR *rs*, the result written to GPR *rd* is 64.

Restrictions:

To be compliant with the MIPS32 and MIPS64 Architecture, software must place the same GPR number in both the *rt* and *rd* fields of the instruction. The operation of the instruction is **UNPREDICTABLE** if the *rt* and *rd* fields of the instruction contain different values.

Operation:

```
temp ← 64
for i in 63.. 0
    if GPR[rs]i = 1 then
        temp ← 63 - i
        break
    endif
endfor
GPR[rd] ← temp
```

Exceptions:

None

31	26	25	16	15	6	5	0
POOL32A 000000	0 0000000000	DERET 1110001101	POOL32AXf 111100				
6	10	10	6				

Format: DERET

EJTAG microMIPS

Purpose: Debug Exception Return

To Return from a debug exception.

Description:

DERET clears execution and instruction hazards, returns from Debug Mode and resumes non-debug execution at the instruction whose address is contained in the *DEPC* register. DERET does not execute the next instruction (i.e. it has no delay slot).

Restrictions:

A DERET placed between an LL and SC instruction does not cause the SC to fail.

If the *DEPC* register with the return address for the DERET was modified by an MTC0 or a DMTC0 instruction, a CP0 hazard exists that must be removed via software insertion of the appropriate number of SSNOP instructions (for implementations of Release 1 of the Architecture) or by an EHB, or other execution hazard clearing instruction (for implementations of Release 2 of the Architecture).

DERET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the DERET returns.

This instruction is legal only if the processor is executing in Debug Mode.

Pre-Release 6: The operation of the processor is **UNDEFINED** if a DERET is executed in the delay slot of a branch or jump instruction. In Release 6, in the absence of delay/forbidden slots, this restriction does not apply.

Operation:


```

DebugDM ← 0
DebugIEXI ← 0
if IsMIPS16Implemented() | (Config3ISA > 0) then
    PC ← DEPC63..1 || 0
    ISAMode ← DEPC0
else
    PC ← DEPC
endif
ClearHazards()

```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: DEXT *rt*, *rs*, *pos*, *size*

microMIPS64

Purpose: Doubleword Extract Bit Field

To extract a bit field from GPR *rs* and store it right-justified into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{ExtractField}(GPR[rs], \text{msbd}, \text{lsb})$

The bit field starting at bit *pos* and extending for *size* bits is extracted from GPR *rs* and stored zero-extended and right-justified in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbd* (the most significant bit of the destination field in GPR *rt*), in instruction bits **15..11**, and *lsb* (least significant bit of the source field in GPR *rs*), in instruction bits **10..6**, as follows:


```
msbd ← size-1
lsb ← pos
msb ← lsb+msbd
```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:

```
0 ≤ pos < 32
0 < size ≤ 32
0 < pos+size ≤ 63
```

Figure 3-3 shows the symbolic operation of the instruction.

Figure 5.5 Operation of the DEXT Instruction



Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* (as derived from *msbd* and *lsb*) and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msbd	lsb	msb	pos	size	Instruction	Comment
$0 \leq \text{msbd} < 32$	$0 \leq \text{lsb} < 32$	$0 \leq \text{msb} < 63$	$0 \leq \text{pos} < 32$	$1 \leq \text{size} \leq 32$	DEXT	The field is 32 bits or less and starts in the right-most word of the doubleword
$0 \leq \text{msbd} < 32$	$32 \leq \text{lsb} < 64$	$32 \leq \text{msb} < 64$	$32 \leq \text{pos} < 64$	$1 \leq \text{size} \leq 32$	DEXTU	The field is 32 bits or less and starts in the left-most word of the doubleword
$32 \leq \text{msbd} < 64$	$0 \leq \text{lsb} < 32$	$32 \leq \text{msb} < 64$	$0 \leq \text{pos} < 32$	$32 < \text{size} \leq 64$	DEXTM	The field is larger than 32 bits and starts in the right-most word of the doubleword

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

Because of the limits on the values of *msbd* and *lsb*, there is no **UNPREDICTABLE** case for this instruction.

Operation:


$$\text{GPR}[\text{rt}] \leftarrow 0^{63-(\text{msbd}+1)} \parallel \text{GPR}[\text{rs}]_{\text{msbd}+\text{lsb}..\text{lsb}}$$

Exceptions:

Reserved Instruction

Programming Notes

The assembler will accept any value of *pos* and *size* that satisfies the relationship $0 < pos+size \leq 64$ and emit DEXT, DEXTM, or DEXTU as appropriate to the values. Programmers should always specify the DEXT mnemonic and let the assembler select the instruction to use.



Format: DEXTM *rt*, *rs*, *pos*, *size*

microMIPS

Purpose: Doubleword Extract Bit Field Middle

To extract a bit field from GPR *rs* and store it right-justified into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{ExtractField}(GPR[rs], msbd, lsb)$

The bit field starting at bit *pos* and extending for *size* bits is extracted from GPR *rs* and stored zero-extended and right-justified in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbminus32* (the most significant bit of the destination field in GPR *rt*, minus 32), in instruction bits 15..11, and *lsb* (least significant bit of the source field in GPR *rs*), in instruction bits 10..6, as follows:

```

msbminus32 ← size-1-32
lsb ← pos
msbd ← msbminus32 + 32
msb ← lsb+msbd

```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:


```

0 ≤ pos < 32
32 < size ≤ 64
32 < pos+size ≤ 64

```

Figure 3-4 shows the symbolic operation of the instruction.

Figure 5.6 Operation of the DEXTM Instruction



Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* (as derived from *msbd* and *lsb*) and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msbd	lsb	msb	pos	size	Instruction	Comment
$0 \leq msbd < 32$	$0 \leq lsb < 32$	$0 \leq msb < 63$	$0 \leq pos < 32$	$1 \leq size \leq 32$	DEXT	The field is 32 bits or less and starts in the right-most word of the doubleword
$0 \leq msbd < 32$	$32 \leq lsb < 64$	$32 \leq msb < 64$	$32 \leq pos < 64$	$1 \leq size \leq 32$	DEXTU	The field is 32 bits or less and starts in the left-most word of the doubleword

msbd	lsb	msb	pos	size	Instruction	Comment
$32 \leq \text{msbd} < 64$	$0 \leq \text{lsb} < 32$	$32 \leq \text{msb} < 64$	$0 \leq \text{pos} < 32$	$32 < \text{size} \leq 64$	DEXTM	The field is larger than 32 bits and starts in the right-most word of the doubleword

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception. The operation is **UNPREDICTABLE** if $(\text{lsb} + \text{msbd} + 1) > 64$.

Operation:

```

msbd ← msbminus32 + 32
if ((lsb + msbd + 1) > 64) then
    UNPREDICTABLE
endif
GPR[rt] ← 063-(msbd+1) || GPR[rs]msbd+lsb..pos


```

Exceptions:

Reserved Instruction

Programming Notes

The assembler will accept any value of *pos* and *size* that satisfies the relationship $0 < \text{pos} + \text{size} \leq 64$ and emit DEXT, DEXTM, or DEXTU as appropriate to the values. Programmers should always specify the DEXT mnemonic and let the assembler select the instruction to use.



Format: DEXTU *rt*, *rs*, *pos*, *size*

microMIPS64

Purpose: Doubleword Extract Bit Field Upper

To extract a bit field from GPR *rs* and store it right-justified into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{ExtractField}(GPR[rs], msbd, lsb)$

The bit field starting at bit *pos* and extending for *size* bits is extracted from GPR *rs* and stored zero-extended and right-justified in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbd* (the most significant bit of the destination field in GPR *rt*), in instruction bits 15..11, and *lsbminus32* (least significant bit of the source field in GPR *rs*, minus 32), in instruction bits 10..6, as follows:

```

msbd ← size-1
lsbminus32 ← pos-32
lsb ← lsbminus32 + 32
msb ← lsb+msbd

```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:


```

32 ≤ pos < 64
0 < size ≤ 32
32 < pos+size ≤ 64

```

Figure 3-5 shows the symbolic operation of the instruction.

Figure 5.7 Operation of the DEXTU Instruction



Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* (as derived from *msbd* and *lsb*) and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msbd	lsb	msb	pos	size	Instruction	Comment
$0 \leq msbd < 32$	$0 \leq lsb < 32$	$0 \leq msb < 63$	$0 \leq pos < 32$	$1 \leq size \leq 32$	DEXT	The field is 32 bits or less and starts in the right-most word of the doubleword
$0 \leq msbd < 32$	$32 \leq lsb < 64$	$32 \leq msb < 64$	$32 \leq pos < 64$	$1 \leq size \leq 32$	DEXTU	The field is 32 bits or less and starts in the left-most word of the doubleword

msbd	lsb	msb	pos	size	Instruction	Comment
$32 \leq \text{msbd} < 64$	$0 \leq \text{lsb} < 32$	$32 \leq \text{msb} < 64$	$0 \leq \text{pos} < 32$	$32 < \text{size} \leq 64$	DEXTM	The field is larger than 32 bits and starts in the right-most word of the doubleword

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception. The operation is **UNPREDICTABLE** if $(\text{lsb} + \text{msbd} + 1) > 64$.

Operation:

```

lsb ← lsbbminus32 + 32
if ((lsb + msbd + 1) > 64) then
    UNPREDICTABLE
endif
GPR[rt] ← 063-(msbd+1) || GPR[rs]msbd+lsb..pos


```

Exceptions:

Reserved Instruction

Programming Notes

The assembler will accept any value of *pos* and *size* that satisfies the relationship $0 < \text{pos} + \text{size} \leq 64$ and emit DEXT, DEXTM, or DEXTU as appropriate to the values. Programmers should always specify the DEXT mnemonic and let the assembler select the instruction to use.



Format: DI
DI rs

microMIPS
microMIPS

Purpose: Disable Interrupts

To return the previous value of the *Status* register and disable interrupts. If DI is specified without an argument, GPR r0 is implied, which discards the previous value of the *Status* register.

Description: $GPR[rs] \leftarrow Status$; $Status_{IE} \leftarrow 0$

The current value of the *Status* register is sign-extended and loaded into general register *rs*. The Interrupt Enable (IE) bit in the *Status* register is then cleared.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

```
data ← Status
GPR[rs] ← sign_extend(data)
StatusIE ← 0
```

Exceptions:


Coprocessor Unusable

Reserved Instruction (Release 1 implementations)

Programming Notes:

The effects of this instruction are identical to those accomplished by the sequence of reading *Status* into a GPR, clearing the IE bit, and writing the result back to *Status*. Unlike the multiple instruction sequence, however, the DI instruction cannot be aborted in the middle by an interrupt or exception.

This instruction creates an execution hazard between the change to the *Status* register and the point where the change to the interrupt enable takes effect. This hazard is cleared by the EHB, JALR.HB, JR.HB, or ERET instructions. Software must not assume that a fixed latency will clear the execution hazard.



Format: DINS *rt*, *rs*, *pos*, *size*

microMIPS64

Purpose: Doubleword Insert Bit Field

To merge a right-justified bit field from GPR *rs* into a specified position in GPR *rt*.

Description: $\text{GPR}[rt] \leftarrow \text{InsertField}(\text{GPR}[rt], \text{GPR}[rs], \text{msb}, \text{lsb})$

The right-most *size* bits from GPR *rs* are merged into the value from GPR *rt* starting at bit position *pos*. The result is placed back in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msb* (the most significant bit of the field), in instruction bits **15..11**, and *lsb* (least significant bit of the field), in instruction bits **10..6**, as follows:

```
msb ← pos+size-1
lsb ← pos
```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:

```
0 ≤ pos < 32
0 < size ≤ 32
0 < pos+size ≤ 32
```

Figure 3-6 shows the symbolic operation of the instruction.

Figure 5.8 Operation of the DINS Instruction

Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msb	lsb	pos	size	Instruction	Comment
$0 \leq \text{msb} < 32$	$0 \leq \text{lsb} < 32$	$0 \leq \text{pos} < 32$	$1 \leq \text{size} \leq 32$	DINS	The field is entirely contained in the right-most word of the doubleword
$32 \leq \text{msb} < 64$	$0 \leq \text{lsb} < 32$	$0 \leq \text{pos} < 32$	$2 \leq \text{size} \leq 64$	DINSM	The field straddles the words of the doubleword
$32 \leq \text{msb} < 64$	$32 \leq \text{lsb} < 64$	$32 \leq \text{pos} < 64$	$1 \leq \text{size} \leq 32$	DINSU	The field is entirely contained in the left-most word of the doubleword

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception. The operation is **UNPREDICTABLE** if *lsb* > *msb*.

Operation:


```
if (lsb > msb) then
    UNPREDICTABLE
endif
GPR[rt] ← GPR[rt]63..msb+1 || GPR[rs]msb-lsb..0 || GPR[rt]lsb-1..0
```

Exceptions:

Reserved Instruction

Programming Notes

The assembler will accept any value of *pos* and *size* that satisfies the relationship $0 < pos+size \leq 64$ and emit DINS, DINSM, or DINSU as appropriate to the values. Programmers should always specify the DINS mnemonic and let the assembler select the instruction to use.



Format: DINSM *rt*, *rs*, *pos*, *size*

microMIPS64

Purpose: Doubleword Insert Bit Field Middle

To merge a right-justified bit field from GPR *rs* into a specified position in GPR *rt*.

Description: $\text{GPR}[rt] \leftarrow \text{InsertField}(\text{GPR}[rt], \text{GPR}[rs], \text{msb}, \text{lsb})$

The right-most *size* bits from GPR *rs* are inserted into the value from GPR *rt* starting at bit position *pos*. The result is placed back in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbminus32* (the most significant bit of the field, minus 32), in instruction bits 15..11, and *lsb* (least significant bit of the field), in instruction bits 10..6, as follows:


```
msbminus32 ← pos+size-1-32
lsb ← pos
msb ← msbminus32 + 32
```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:

```
0 ≤ pos < 32
2 ≤ size ≤ 64
32 < pos+size ≤ 64
```

Figure 3-7 shows the symbolic operation of the instruction.

Figure 5.9 Operation of the DINSM Instruction



Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msb	lsb	pos	size	Instruction	Comment
$0 \leq msb < 32$	$0 \leq lsb < 32$	$0 \leq pos < 32$	$1 \leq size \leq 32$	DINS	The field is entirely contained in the right-most word of the doubleword
$32 \leq msb < 64$	$0 \leq lsb < 32$	$0 \leq pos < 32$	$2 \leq size \leq 64$	DINSM	The field straddles the words of the doubleword
$32 \leq msb < 64$	$32 \leq lsb < 64$	$32 \leq pos < 64$	$1 \leq size \leq 32$	DINSU	The field is entirely contained in the left-most word of the doubleword

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception. Because of the instruction format, *lsb* can never be greater than *msb*, so there is no **UNPREDICATABLE** case for this instruction.

Operation:


$$msb \leftarrow msb_{\text{minus}32} + 32$$

$$GPR[rt] \leftarrow GPR[rt]_{63..msb+1} \parallel GPR[rs]_{msb-lsb..0} \parallel GPR[rt]_{lsb-1..0}$$
Exceptions:

Reserved Instruction

Programming Notes

The assembler will accept any value of *pos* and *size* that satisfies the relationship $0 < pos+size \leq 64$ and emit DINS, DINSM, or DINSU as appropriate to the values. Programmers should always specify the DINS mnemonic and let the assembler select the instruction to use.



Format: DINSU *rt*, *rs*, *pos*, *size*

microMIPS64

Purpose: Doubleword Insert Bit Field Upper

To merge a right-justified bit field from GPR *rs* into a specified position in GPR *rt*.

Description: $\text{GPR}[rt] \leftarrow \text{InsertField}(\text{GPR}[rt], \text{GPR}[rs], \text{msb}, \text{lsb})$

The right-most *size* bits from GPR *rs* are inserted into the value from GPR *rt* starting at bit position *pos*. The result is placed back in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbminus32* (the most significant bit of the field, minus 32), in instruction bits 15..11, and *lsbminus32* (least significant bit of the field, minus 32), in instruction bits 10..6, as follows:

```

msbminus32 ← pos+size-1-32
lsbminus32 ← pos-32
msb ← msbminus32 + 32
lsb ← lsbminus32 + 32

```

For this instruction, the values of *pos* and *size* must satisfy all of the following relations:

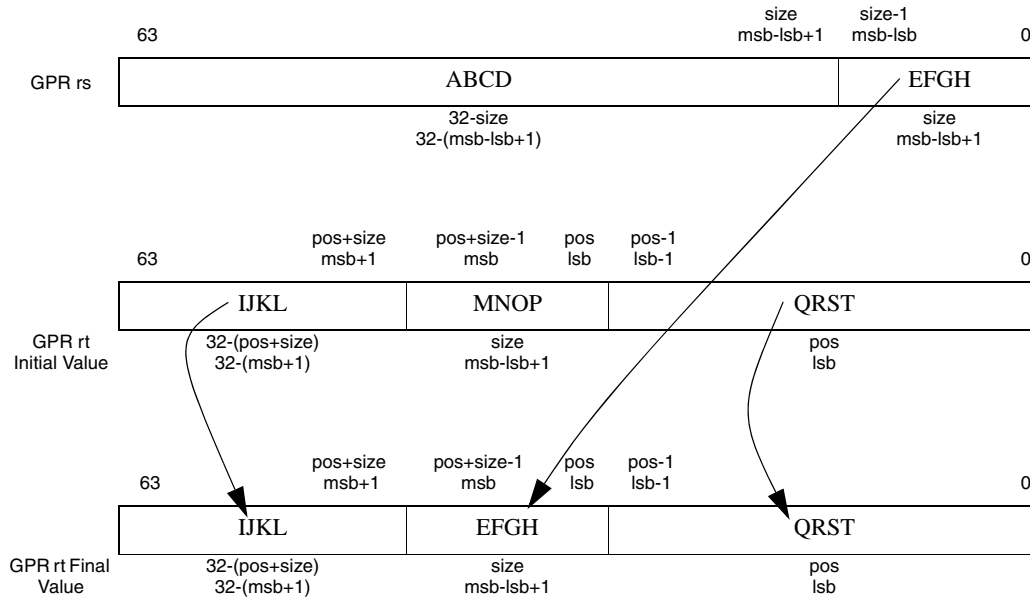
```

32 ≤ pos < 64
1 ≤ size ≤ 32
32 < pos+size ≤ 64

```

Figure 3-8 shows the symbolic operation of the instruction.

Figure 5.10 Operation of the DINSU Instruction



Three instructions are required to access any legal bit field within the doubleword, as a function of the *msb* and *lsb* of the field (which implies restrictions on *pos* and *size*), as follows:

msb	lsb	pos	size	Instruction	Comment
$0 \leq msb < 32$	$0 \leq lsb < 32$	$0 \leq pos < 32$	$1 \leq size \leq 32$	DINS	The field is entirely contained in the right-most word of the doubleword
$32 \leq msb < 64$	$0 \leq lsb < 32$	$0 \leq pos < 32$	$2 \leq size \leq 64$	DINSM	The field straddles the words of the doubleword
$32 \leq msb < 64$	$32 \leq lsb < 64$	$32 \leq pos < 64$	$1 \leq size \leq 32$	DINSU	The field is entirely contained in the left-most word of the doubleword

Restrictions:

In implementations pre-Release 2 of the architecture, the instruction resulted in a Reserved Instruction exception.

The operation is **UNPREDICTABLE** if $lsb > msb$.

Operation:

```

lsb ← lsbminus32 + 32
msb ← msbminus32 + 32
if (lsb > msb) then
    UNPREDICTABLE
endif
GPR[rt] ← GPR[rt]63..msb+1 || GPR[rs]msb-lsb..0 || GPR[rt]lsb-1..0


```

Exceptions:

Reserved Instruction

Programming Notes

The assembler accepts any value of *pos* and *size* that satisfies the relationship $0 < pos+size \leq 64$ and emit DINS, DINSM, or DINSU as appropriate to the values. Programmers should always specify the DINS mnemonic and let the assembler select the instruction to use.



Format: DIV.fmt
 DIV.S fd, fs, ft
 DIV.D fd, fs, ft

microMIPS
 microMIPS

Purpose: Floating Point Divide

To divide FP values.

Description: $FPR[fd] \leftarrow FPR[fs] / FPR[ft]$

The value in FPR *fs* is divided by the value in FPR *ft*. The result is calculated to infinite precision, rounded according to the current rounding mode in *FCSR*, and placed into FPR *fd*. The operands and result are values in format *fmt*.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

Operation:

StoreFPR (fd, fmt, ValueFPR(fs, fmt) / ValueFPR(ft, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Invalid Operation, Unimplemented Operation, Division-by-zero, Overflow, Underflow

31	26	25	21	20	16	15	11	10	9	0	
POOL32A 000000	rt	rs	rd	0	DIV 0100011000						
POOL32A 000000	rt	rs	rd	0	MOD 0101011000						
POOL32A 000000	rt	rs	rd	0	DIVU 0110011000						
POOL32A 000000	rt	rs	rd	0	MODU 0111011000						
6	5	5	5	1	10						
31	26	25	21	20	16	15	11	10	9	8	0
POOL32S 010110	rt	rs	rd	00	DDIV 100011000						
POOL32S 010110	rt	rs	rd	00	DMOD 101011000						
POOL32S 010110	rt	rs	rd	00	DDIVU 110011000						
POOL32S 010110	rt	rs	rd	00	DMODU 111011000						
6	5	5	5	2	9						

Format: DIV MOD DIVU MODU DDIV DMOD DDIVU DMODU

DIV rd,rs,rt
 MOD rd,rs,rt
 DIVU rd,rs,rt
 MODU rd,rs,rt
 DDIV rd,rs,rt
 DMOD rd,rs,rt
 DDIVU rd,rs,rt
 DMODU rd,rs,rt

microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6

Purpose: Divide Integers (with result to GPR)

DIV: Divide Words Signed

MOD: Modulo Words Signed

DIVU: Divide Words Unsigned

MODU: Modulo Words Unsigned

DDIV: Divide Doublewords Signed

DMOD: Modulo Doublewords Signed

DDIVU: Divide Doublewords Unsigned

DMODU: Modulo Doublewords Unsigned

Description:

DIV: GPR[rd] ← sign_extend.32(divide.signed(GPR[rs], GPR[rt]))
 MOD: GPR[rd] ← sign_extend.32(modulo.signed(GPR[rs], GPR[rt]))
 DIVU: GPR[rd] ← sign_extend.32(divide.unsigned(GPR[rs], GPR[rt]))
 MODU: GPR[rd] ← sign_extend.32(modulo.unsigned(GPR[rs], GPR[rt]))

```

DDIV: GPR[rd] ← divide.signed( GPR[rs], GPR[rt] )
DMOD: GPR[rd] ← modulo.signed( GPR[rs], GPR[rt] )
DDIVU: GPR[rd] ← divide.unsigned( GPR[rs], GPR[rt] )
DMODU: GPR[rd] ← modulo.unsigned( GPR[rs], GPR[rt] )

```

The Release 6 divide and modulo instructions divide the operands in GPR *rs* and GPR *rt*, and place the quotient or remainder in GPR *rd*.

For each of the div/mod operator pairs DIV/MOD, DIVU/MODU, DDIV/DMOD, DDIVU/DMODU the results satisfy the equation $(A \text{ div } B) * B + (A \text{ mod } B) = A$, where $(A \text{ mod } B)$ has same sign as the dividend *A*, and $\text{abs}(A \text{ mod } B) < \text{abs}(B)$. This equation uniquely defines the results.

NOTE: if the divisor $B=0$, this equation cannot be satisfied, and the result is UNPREDICTABLE. This is commonly called “truncated division”.

DIV performs a signed 32-bit integer division, and places the 32-bit quotient result in the destination register.

MOD performs a signed 32-bit integer division, and places the 32-bit remainder result in the destination register. The remainder result has the same sign as the dividend.

DIVU performs an unsigned 32-bit integer division, and places the 32-bit quotient result in the destination register.

MODU performs an unsigned 32-bit integer division, and places the 32-bit remainder result in the destination register.

DDIV performs a signed 64-bit integer division, and places the 64-bit quotient result in the destination register.

DMOD performs a signed 64-bit integer division, and places the 64-bit remainder result in the destination register. The remainder result has the same sign as the dividend.

DDIVU performs an unsigned 64-bit integer division, and places the 64-bit quotient result in the destination register.

DMODU performs an unsigned 64-bit integer division, and places the 64-bit remainder result in the destination register.

Restrictions:

If the divisor in GPR *rt* is zero, the result value is UNPREDICTABLE.

On a 64-bit CPU, the 32-bit signed divide (DIV) and modulo (MOD) instructions are UNPREDICTABLE if inputs are not signed extended 32-bit integers.

Special provision is made for the inputs to unsigned 32-bit divide and modulo on a 64-bit CPU. Since many 32-bit instructions sign extend 32 bits to 64 even for unsigned computation, properly sign extended numbers must be accepted as input, and truncated to 32 bits, clearing bits 32-63. However, it is also desirable to accept zero extended 32-bit integers, with bits 32-63 all 0.

On a 64-bit CPU, DIVU and MODU are UNPREDICTABLE if their inputs are not zero or sign extended 32-bit integers.

On a 64-bit CPU, the 32-bit divide and modulo instructions, both signed and unsigned, sign extend the result as if it is a 32-bit signed integer.

DDIV, DMOD, DDIVU, DMODU: Reserved Instruction exception if 64-bit instructions are not enabled.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Release 6 divide instructions have the same opcode mnemonic as the pre-Release 6 divide instructions (DIV, DIVU, DDIV, DDIVU). The instruction encodings are different, as are the instruction semantics: the Release 6 instruction produces only the quotient, whereas the pre-Release 6 instruction produces quotient and remainder in HI/LO registers respectively, and separate modulo instructions are required to obtain the remainder.

The assembly syntax distinguishes the Release 6 from the pre-Release 6 divide instructions. For example, Release 6

“DIV rd,rs,rt” specifies 3 register operands, versus pre-Release 6 “DIV rs,rt”, which has only two register arguments, with the HI/LO registers implied. Some assemblers accept the pseudo-instruction syntax “DIV rd,rs,rt” and expand it to do “DIV rs,rt;MFHI rd”. Phrases such as “DIV with GPR output” and “DIV with HI/LO output” may be used when disambiguation is necessary.

Pre-Release 6 divide instructions that produce quotient and remainder in the HI/LO registers produce a Reserved Instruction exception on Release 6. In the future, the instruction encoding may be reused for other instructions.

Programming Notes:

Because the divide and modulo instructions are defined to not trap if dividing by zero, it is safe to emit code that checks for zero-divide after the divide or modulo instruction.

Operation

```
DDIV, DMOD, DDIVU, DMODU:
    if not Are64bitOperationsEnabled then SignalException(ReservedInstruction) endif
    if NotWordValue(GPR[rs]) then UNPREDICTABLE endif
    if NotWordValue(GPR[rt]) then UNPREDICTABLE endif

/* recommended implementation: ignore bits 32-63 for DIV, MOD, DIVU, MODU */

DIV, MOD:
    s1 ← signed_word(GPR[rs])
    s2 ← signed_word(GPR[rt])
DIVU, MODU:
    s1 ← unsigned_word(GPR[rs])
    s2 ← unsigned_word(GPR[rt])
DDIV, DMOD:
    s1 ← signed_doubleword(GPR[rs])
    s2 ← signed_doubleword(GPR[rt])
DDIVU, DMODU:
    s1 ← unsigned_doubleword(GPR[rs])
    s2 ← unsigned_doubleword(GPR[rt])

DIV, DIVU, DDIV, DDIVU:
    quotient ← s1 div s2
MOD, MODU, DMOD, DMODU:
    remainder ← s1 mod s2

DIV:  GPR[rd] ← sign_extend.32( quotient )
MOD:  GPR[rd] ← sign_extend.32( remainder )
DIVU: GPR[rd] ← sign_extend.32( quotient )
MODU: GPR[rd] ← sign_extend.32( remainder )
DDIV: GPR[rd] ← quotient
DMOD: GPR[rd] ← remainder
DDIVU: GPR[rd] ← quotient
DMODU: GPR[rd] ← remainder
/* end of instruction */


where

function zero_or_sign_extended.32(val)
    if value63..32 = (value31)32 then return true
    if value63..32 = (0)32 then return true
    return false
end function
```

Exceptions:

DIV, MOD, DIVU, MODU: No arithmetic exceptions occur. Division by zero produces an UNPREDICTABLE result.

DDIV, DMOD, DDIVU, DMODU: Reserved Instruction.



Format: DMFC0 *rt*, *rs*
DMFC0 *rt*, *rs*, *sel*

microMIPS64
microMIPS64

Purpose: Doubleword Move from Coprocessor 0

To move the contents of a coprocessor 0 register to a general purpose register (GPR).

Description: $GPR[rt] \leftarrow CPR[0,rs,sel]$

The contents of the coprocessor 0 register are loaded into GPR *rt*. Note that not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be zero.

Restrictions:


The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rs* and *sel*, or if the coprocessor 0 register specified by *rd* and *sel* is a 32-bit register.

Operation:

```
// 'Width' returns width (32/64) of data returned by CPR
if ((Width(CPR[0,rs,sel]) = 32) and (ConfigAR>=2) ) then
    dataword ← CPR[0,rs,sel]
    GPR[rt] ← { x0000_0000 || dataword }
elseif ((Width(CPR[0,rs,sel]) = 32) and (ConfigAR<2)) then
    UNDEFINED
else
    datadoubleword ← CPR[0,rs,sel]
    GPR[rt] ← datadoubleword
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: DMFC1 *rt*, *fs*

microMIPS64

Purpose: Doubleword Move from Floating Point

To move a doubleword from an FPR to a GPR.

Description: $GPR[rt] \leftarrow FPR[fs]$

The contents of FPR *fs* are loaded into GPR *rt*..

Restrictions:

Operation:

```
datadoubleword ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)
GPR[rt] ← datadoubleword
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Historical Information:

For MIPS III, the contents of GPR *rt* are undefined for the instruction immediately following DMFC1.

31	26	25	21	20	16	15	6	5	0
POOL32A 000000			rt		Impl		DMFC2 0110110100		POOL32Axf 111100
6			5		5		10		6

Format: DMFC2 rt, rd
DMFC2, rt, rd, sel

microMIPS64
microMIPS64

The syntax shown above is an example using DMFC1 as a model. The specific syntax is implementation dependent.

Purpose: Doubleword Move from Coprocessor 2

To move a doubleword from a coprocessor 2 register to a GPR.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{CP2CPR}[\text{Impl}]$

The contents of the coprocessor 2 register denoted by the *impl* field is loaded into GPR *rt*. The interpretation of the *impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The results are **UNPREDICTABLE** if *impl* specifies a coprocessor 2 register that does not exist, or if the coprocessor 2 register specified by *rd* and *sel* is a 32-bit register.

Operation:

```
datadoubleword ← CP2CPR[Impl]
GPR[rt] ← datadoubleword
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: DMTC0 *rt*, *rs*
DMTC0 *rt*, *rs*, *sel*

microMIPS64
microMIPS64

Purpose: Doubleword Move to Coprocessor 0

To move a doubleword from a GPR to a coprocessor 0 register.

Description: $CPR[0,rs,sel] \leftarrow GPR[rt]$

The contents of GPR *rt* are loaded into the coprocessor 0 register specified in the *rd* and *sel* fields. Not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be zero.

Restrictions:

The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rs* and *sel*, or if the coprocessor 0 register specified by *rd* and *sel* is a 32-bit register.

Operation:

```
// 'Width' returns width (32/64) of data returned by CPR.
if ((Width(CPR[0,rs,sel]) = 32) and (ConfigAR >= 2) ) then
    dataword <- GPR[rt]31:0
    CPR[0,rs,sel] <- dataword
elseif ((Width(CPR[0,rs,sel]) = 32) and (ConfigAR < 2)) then
    UNDEFINED
else
    datadoubleword <- GPR[rt]
    CPR[0,rs,sel] <- datadoubleword
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

31	26	25	21	20	16	15	14	13	6	5	0	
POOL32F 010101			rt		fs		0 00	DMTC1 10110000			POOL32Fxf 111011	
6			5		5		2	8			6	

Format: DMTC1 *rt*, *fs*

microMIPS64

Purpose: Doubleword Move to Floating Point

To copy a doubleword from a GPR to an FPR.

Description: $\text{FPR}[\text{fs}] \leftarrow \text{GPR}[\text{rt}]$

The doubleword contents of GPR *rt* are placed into FPR *fs*.

Restrictions:

Operation:


```
datadoubleword ← GPR[rt]
StoreFPR(fs, UNINTERPRETED_DOUBLEWORD, datadoubleword)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Historical Information:

For MIPS III, the contents of FPR *fs* are undefined for the instruction immediately following DMTC1.



Format: DMTC2 *rt*, *Impl*
DMTC2 *rt*, *Impl*, *sel*

microMIPS64
microMIPS64

The syntax shown above is an example using DMTC1 as a model. The specific syntax is implementation dependent.

Purpose: Doubleword Move to Coprocessor 2

To move a doubleword from a GPR to a coprocessor 2 register.

Description: $CP2CPR[Impl] \leftarrow GPR[rt]$

The contents GPR *rt* are loaded into the coprocessor 2 register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The results are **UNPREDICTABLE** if *Impl* specifies a coprocessor 2 register that does not exist, or if the coprocessor 2 register specified by *rd* and *sel* is a 32-bit register.

Operation:

$datadoubleword \leftarrow GPR[rt]$
 $CP2CPR[Impl] \leftarrow datadoubleword$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: DROTR *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Rotate Right

To execute a logical right-rotate of a doubleword by a fixed amount—0 to 31 bits.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \times (\text{right}) \text{ sa}$

The doubleword contents of GPR *rs* are rotated right; the result is placed in GPR *rt*. The bit-rotate amount in the range 0 to 31 is specified by *sa*.


Restrictions:

Operation:

$$\begin{aligned} \text{s} &\leftarrow 0 \mid \mid \text{sa} \\ \text{GPR}[\text{rt}] &\leftarrow \text{GPR}[\text{rs}]_{\text{s}-1..0} \mid \mid \text{GPR}[\text{rs}]_{63..\text{s}} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DROTR32 *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Rotate Right Plus 32

To execute a logical right-rotate of a doubleword by a fixed amount—32 to 63 bits

Description: $GPR[rt] \leftarrow GPR[rs] \times (\text{right}) (saminus32+32)$

The 64-bit doubleword contents of GPR *rs* are rotated right; the result is placed in GPR *rt*. The bit-rotate amount in the range 32 to 63 is specified by *saminus32+32*.

Restrictions:

Operation:


```

s      ← 1 || sa      /* 32+saminus32 */
GPR[rt] ← GPR[rs]s-1..0 || GPR[rt]63..s

```

Exceptions:

Reserved Instruction



Format: DROTRV rd, rt, rs

microMIPS64

Purpose: Doubleword Rotate Right Variable

To execute a logical right-rotate of a doubleword by a variable number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \times (\text{right}) GPR[rs]$

The 64-bit doubleword contents of GPR *rt* are rotated right; the result is placed in GPR *rd*. The bit-rotate amount in the range 0 to 63 is specified by the low-order 6 bits in GPR *rs*.


Restrictions:

Operation:

$$\begin{aligned} s &\leftarrow GPR[rs]_{5..0} \\ GPR[rd] &\leftarrow GPR[rt]_{s-1..0} \parallel GPR[rt]_{63..s} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSBH *rt*, *rs*

microMIPS64

Purpose: Doubleword Swap Bytes Within Halfwords

To swap the bytes within each halfword of GPR *rs* and store the value into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{SwapBytesWithinHalfwords}(GPR[rs])$

Within each halfword of GPR *rs* the bytes are swapped and stored in GPR *rt*.

Restrictions:

In implementations Release 1 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

$$GPR[rt] \leftarrow \begin{matrix} GPR[s]_{55..48} & || & GPR[s]_{63..56} & || & GPR[s]_{39..32} & || & GPR[s]_{47..40} & || \\ GPR[s]_{23..16} & || & GPR[s]_{31..24} & || & GPR[s]_{7..0} & || & GPR[s]_{15..8} \end{matrix}$$


Exceptions:

Reserved Instruction

Programming Notes:

The DSBH and DSHD instructions can be used to convert doubleword data of one endianness to the other endianness. For example:

```
ld    t0, 0(a1)      /* Read doubleword value */
dsbh  t0, t0          /* Convert endiannes of the halfwords */
dshd  t0, t0          /* Swap the halfwords within the doublewords */
```



Format: DSHD *rt*, *rs*

microMIPS64

Purpose: Doubleword Swap Halfwords Within Doublewords

To swap the halfwords of GPR *rs* and store the value into GPR *rt*.

Description: $\text{GPR}[rt] \leftarrow \text{SwapHalfwordsWithinDoublewords}(\text{GPR}[rs])$

The halfwords of GPR *rs* are swapped and stored in GPR *rt*.

Restrictions:

In implementations of Release 1 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

$$\geq \text{GPR}[rt] \leftarrow \text{GPR}[rs]_{15..0} \parallel \text{GPR}[rs]_{31..16} \parallel \text{GPR}[rs]_{47..32} \parallel \text{GPR}[rs]_{63..48}$$

Exceptions:


Reserved Instruction

Programming Notes:

The DSBH and DSHD instructions can be used to convert doubleword data of one endianness to the other endianness.

For example:

```
ld      t0, 0(a1)      /* Read doubleword value */
dsbh    t0, t0          /* Convert endiannes of the halfwords */
dshd    t0, t0          /* Swap the halfwords within the doublewords */
```



Format: DSLL *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Left Logical

To execute a left-shift of a doubleword by a fixed amount—0 to 31 bits

Description: $GPR[rt] \leftarrow GPR[rs] \ll sa$

The 64-bit doubleword contents of GPR *rs* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 0 to 31 is specified by *sa*.

Restrictions:


Operation:

$$s \leftarrow 0 \parallel sa$$

$$GPR[rt] \leftarrow GPR[rs]_{(63-s) \dots 0} \parallel 0^s$$

Exceptions:

Reserved Instruction



Format: DSLL32 *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Left Logical Plus 32

To execute a left-shift of a doubleword by a fixed amount—32 to 63 bits

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \ll (\text{sa} + 32)$

The 64-bit doubleword contents of GPR *rs* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 0 to 31 is specified by *sa*.

Restrictions:

Operation:

$$\begin{aligned} s &\leftarrow 1 \mid \mid sa \quad /* \ 32+sa \ */ \\ \text{GPR}[\text{rt}] &\leftarrow \text{GPR}[\text{rs}]_{(63-s) \dots 0} \mid \mid 0^s \end{aligned}$$

Exceptions:

Reserved Instruction

31	26	25	21	20	16	15	11	10	9	8	0	
POOL32S 010110			rt		rs		rd		0 00		DSLLV 000010000	
6			5		5		5		2		9	

Format: DSLLV rd, rt, rs

microMIPS64

Purpose: Doubleword Shift Left Logical Variable

To execute a left-shift of a doubleword by a variable number of bits.

Description: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rt}] \ll \text{GPR}[\text{rs}]$

The 64-bit doubleword contents of GPR *rt* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit-shift amount in the range 0 to 63 is specified by the low-order 6 bits in GPR *rs*.


Restrictions:

Operation:

$$\begin{aligned} s &\leftarrow \text{GPR}[\text{rs}]_{5..0} \\ \text{GPR}[\text{rd}] &\leftarrow \text{GPR}[\text{rt}]_{(63-s) .. 0} \parallel 0^s \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSRA *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Right Arithmetic

To execute an arithmetic right-shift of a doubleword by a fixed amount—0 to 31 bits.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \gg \text{sa}$ (arithmetic)

The 64-bit doubleword contents of GPR *rs* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 0 to 31 is specified by *sa*.


Restrictions:

Operation:

$$\begin{aligned} s &\leftarrow 0 \mid \mid sa \\ \text{GPR}[\text{rt}] &\leftarrow (\text{GPR}[\text{rs}]_{63})^s \mid \mid \text{GPR}[\text{rs}]_{63..s} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSRA32 *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Right Arithmetic Plus 32

To execute an arithmetic right-shift of a doubleword by a fixed amount—32 to 63 bits

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \gg (\text{sa}+32)$ (arithmetic)

The doubleword contents of GPR *rs* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 32 to 63 is specified by *sa*+32.


Restrictions:

Operation:

$$\begin{aligned} \text{s} &\leftarrow 1 \mid \mid \text{sa} \quad /* \text{ 32+sa */} \\ \text{GPR}[\text{rt}] &\leftarrow (\text{GPR}[\text{rs}]_{63})^{\text{s}} \mid \mid \text{GPR}[\text{rs}]_{63..\text{s}} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSRAV rd, rt, rs

microMIPS64

Purpose: Doubleword Shift Right Arithmetic Variable

To execute an arithmetic right-shift of a doubleword by a variable number of bits.

Description: $GPR[rd] \leftarrow GPR[rt] \gg GPR[rs]$ (arithmetic)

The doubleword contents of GPR *rt* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rd*. The bit-shift amount in the range 0 to 63 is specified by the low-order 6 bits in GPR *rs*.


Restrictions:

Operation:

$$GPR[rd] \leftarrow (GPR[rt]_{63})^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRL *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Right Logical

To execute a logical right-shift of a doubleword by a fixed amount—0 to 31 bits.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rs}] \gg \text{sa}$ (logical)

The doubleword contents of GPR *rs* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 0 to 31 is specified by *sa*.


Restrictions:

Operation:

$$\begin{aligned} \text{s} &\leftarrow 0 \mid \mid \text{sa} \\ \text{GPR}[\text{rt}] &\leftarrow 0^{\text{s}} \mid \mid \text{GPR}[\text{rs}]_{63..\text{s}} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSRL32 *rt*, *rs*, *sa*

microMIPS64

Purpose: Doubleword Shift Right Logical Plus 32

To execute a logical right-shift of a doubleword by a fixed amount—32 to 63 bits

Description: $GPR[rt] \leftarrow GPR[rs] \gg (saminus32+32)$ (logical)

The 64-bit doubleword contents of GPR *rs* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rt*. The bit-shift amount in the range 32 to 63 is specified by *saminus32+32*.

Restrictions:

Operation:

```

s      ← 1 || sa /* 32+saminus32 */
GPR[rt] ← 0s || GPR[rs]63..s

```

Exceptions:

Reserved Instruction

31	26	25	21	20	16	15	11	10	9	8	0	
POOL32S 010110			rt		rs		rd		0 00		DSRLV 001010000	
6			5		5		5		2		9	

Format: DSRLV rd, rt, rs

microMIPS64

Purpose: Doubleword Shift Right Logical Variable

To execute a logical right-shift of a doubleword by a variable number of bits

Description: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rt}] \gg \text{GPR}[\text{rs}]$ (logical)

The 64-bit doubleword contents of GPR *rt* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit-shift amount in the range 0 to 63 is specified by the low-order 6 bits in GPR *rs*.


Restrictions:

Operation:

$$\begin{aligned} s &\leftarrow \text{GPR}[\text{rs}]_{5..0} \\ \text{GPR}[\text{rd}] &\leftarrow 0^s \mid \mid \text{GPR}[\text{rt}]_{63..s} \end{aligned}$$

Exceptions:

Reserved Instruction



Format: DSUB rd, rs, rt

microMIPS64

Purpose: Doubleword Subtract

To subtract 64-bit integers; trap on overflow

Description: $GPR[rd] \leftarrow GPR[rs] - GPR[rt]$

The 64-bit doubleword value in GPR *rt* is subtracted from the 64-bit value in GPR *rs* to produce a 64-bit result. If the subtraction results in 64-bit 2's complement arithmetic overflow, then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 64-bit result is placed into GPR *rd*.

Restrictions:

Operation:

```
temp ← (GPR[rs]63 || GPR[rs]) - (GPR[rt]63 || GPR[rt])
if (temp64 ≠ temp63) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp63..0
endif
```

Exceptions:

Integer Overflow, Reserved Instruction

Programming Notes:

DSUBU performs the same arithmetic operation but does not trap on overflow.

31	26	25	21	20	16	15	11	10	9	8	0	
POOL32S 010110			rt		rs		rd		0 00		DSUBU 111010000	
6			5		5		5		2		9	

Format: DSUBU rd, rs, rt

microMIPS64

Purpose: Doubleword Subtract Unsigned

To subtract 64-bit integers

Description: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}] - \text{GPR}[\text{rt}]$

The 64-bit doubleword value in GPR *rt* is subtracted from the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

Operation: 64-bit processors

$\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}] - \text{GPR}[\text{rt}]$

Exceptions:

Reserved Instruction

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.

31	26	25	21	20	16	15	6	5	0
POOL32A 000000	rt		0 00000	DVP 0001100101			POOL32AXf 111100		
6	5		5	10			6		

Format: DVP rt

microMIPS Release 6

Purpose: Disable Virtual Processor

To disable all virtual processors in a physical core other than the virtual processor that issued the instruction.

Description: $GPR[rt] \leftarrow VPControl$; $VPControl_{DIS} \leftarrow 1$

Disabling a virtual processor means that instruction fetch is terminated, and all outstanding instructions for the affected virtual processor(s) must be complete before the DVP itself is allowed to retire. Any outstanding events such as hardware instruction or data prefetch, or page-table walks must also be terminated.

The DVP instruction has implicit SYNC(*stype*=0) semantics but with respect to the other virtual processors in the physical core.

After all other virtual processors have been disabled, $VPControl_{DIS}$ is set. Prior to modification and if *rt* is non-zero, sign-extended $VPControl$ is written to $GPR[rt]$. If DVP is specified without *rt*, then *rt* must be 0.

DVP may also take effect on a virtual processor that has executed a WAIT or a PAUSE instruction. If a virtual processor has executed a WAIT instruction, then it cannot resume execution on an interrupt until an EVP has been executed. If the EVP is executed before the interrupt arrives, then the virtual processor resumes in a state as if the DVP had not been executed, that is, it waits for the interrupt.

If a virtual processor has executed a PAUSE instruction, then it cannot resume execution until an EVP has been executed, even if LLbit is cleared. If an EVP is executed before the LLbit is cleared, then the virtual processor resumes in a state as if the DVP has not been executed, that is, it waits for the LLbit to clear.

The execution of a DVP must be followed by the execution of an EVP. The execution of an EVP causes execution to resume immediately—where applicable—on all other virtual processors, as if the DVP had not been executed. The execution is completely restorable after the EVP. If an event occurs in between the DVP and EVP that renders state of the virtual processor UNPREDICTABLE (such as power-gating), then the effect of EVP is UNPREDICTABLE.

DVP may only take effect if $VPControl_{DIS}=0$. Otherwise it is treated as a NOP instruction.

If a virtual processor is disabled due to a DVP, then interrupts are also disabled for the virtual processor, that is, logically $Status_{IE}=0$. $Status_{IE}$ for the target virtual processors though is not cleared though as software cannot access state on the virtual processors that have been disabled. Similarly, deferred exceptions will not cause a disabled virtual processor to be re-enabled for execution, at least until execution is re-enabled by the EVP instruction. The virtual processor that executes the DVP, however, continues to be interruptible.

In an implementation, the ability of a virtual processor to execute instructions may also be under control external to the physical core which contains the virtual processor. If disabled by DVP, a virtual processor must not resume fetch in response to the assertion of this external signal to enable fetch. Conversely, if fetch is disabled by such external control, then execution of EVP will not cause fetch to resume at a target virtual processor for which the control is deasserted.

This instruction never executes speculatively. It must be the oldest unretired instruction to take effect.

This instruction is only available in Release 6 implementations. For implementations that do not support multi-threading ($Config5_{VP}=0$), this instruction must be treated as a NOP instruction.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 6 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

The pseudo-code below assumes that the DVP is executed by virtual processor 0, while the target virtual processor is numbered 'n', where n is each of all remaining virtual processors.

```

if (VPControlDIS = 0)

    // Pseudo-code in italics provides recommended action wrt other VPs
    disable_fetch(VPn) {
        if PAUSE(VPn) retires prior or at disable event
            then VPn execution is not resumed if LLbit is cleared prior to EVP
    }
    disable_interrupt(VPn) {
        if WAIT(VPn) retires prior or at disable event
            then interrupts are ignored by VPn until EVP
    }
    // DVP0 not retired until instructions for VPn completed
    while (VPn outstanding instruction)
        DVP0 unretired
    endwhile

endif

data ← VPControl
GPR[rt] ← sign_extend(data)
VPControlDIS ← 1

```

Exceptions:

Coprocessor Unusable

Reserved Instruction (pre-Release 6 implementations)

Programming Notes:

DVP may disable execution in the target virtual processor regardless of the operating mode - kernel, supervisor, user. Kernel software may also be in a critical region, or in a high-priority interrupt handler when the disable occurs. Since the instruction is itself privileged, such events are considered acceptable.

Before executing an EVP in a DVP/EVP pair, software should first read VPControl_{DIS}, returned by DVP, to determine whether the virtual processors are already disabled. If so, the DVP/EVP sequence should be abandoned. This step allows software to safely nest DVP/EVP pairs.

Privileged software may use DVP/EVP to disable virtual processors on a core, such as for the purpose of doing a cache flush without interference from other processes in a system with multiple virtual processors or physical cores.


DVP (and EVP) may be used in other cases such as for power-savings or changing state that is applicable to all virtual processors in a core, such as virtual processor scheduling priority, as described below:

```

ll t0 0(a0)
dvp    // disable all other virtual processors
pause  // wait for LLbit to clear
evp    // enable all other virtual processors

```

```
ll t0 0(a0)
dvp    // disable all other virtual processors
<change core-wide state>
evp    // enable all other virtual processors
```

Format: EHB

microMIPS

Purpose: Execution Hazard Barrier

To stop instruction execution until all execution hazards have been cleared.

Description:

EHB is used to denote execution hazard barrier. The actual instruction is interpreted by the hardware as SLL r0, r0, 3.

This instruction alters the instruction issue behavior on a pipelined processor by stopping execution until all execution hazards have been cleared. Other than those that might be created as a consequence of setting *Status_{CU0}*, there are no execution hazards visible to an unprivileged program running in User Mode. All execution hazards created by previous instructions are cleared for instructions executed immediately following the EHB. The EHB instruction does not clear instruction hazards—such hazards are cleared by the JALR.HB, JR.HB, and ERET instructions.

Restrictions:

None

Operation:


`ClearExecutionHazards()`

Exceptions:

None

Programming Notes:

In Release 2 implementations, this instruction resolves all execution hazards. On a superscalar processor, EHB alters the instruction issue behavior in a manner identical to SSNOP. For backward compatibility with Release 1 implementations, the last of a sequence of SSNOPs can be replaced by an EHB. In Release 1 implementations, the EHB will be treated as an SSNOP, thereby preserving the semantics of the sequence. In Release 2 implementations, replacing the final SSNOP with an EHB should have no performance effect because a properly sized sequence of SSNOPs will have already cleared the hazard. As EHB becomes the standard in MIPS implementations, the previous SSNOPs can be removed, leaving only the EHB.



Format: EI
EI rs

microMIPS
microMIPS

Purpose: Enable Interrupts

To return the previous value of the *Status* register and enable interrupts. If EI is specified without an argument, GPR r0 is implied, which discards the previous value of the *Status* register.

Description: $GPR[rt] \leftarrow Status$; $Status_{IE} \leftarrow 1$

The current value of the *Status* register is sign-extended and loaded into general register *rt*. The Interrupt Enable (*IE*) bit in the *Status* register is then set.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

```
data ← Status
GPR[rs] ← sign_extend(data)
StatusIE ← 1
```

Exceptions:

Coprocessor Unusable

Reserved Instruction (Release 1 implementations)

Programming Notes:

The effects of this instruction are identical to those accomplished by the sequence of reading *Status* into a GPR, setting the *IE* bit, and writing the result back to *Status*. Unlike the multiple instruction sequence, however, the EI instruction cannot be aborted in the middle by an interrupt or exception.

This instruction creates an execution hazard between the change to the Status register and the point where the change to the interrupt enable takes effect. This hazard is cleared by the EHB, JALR.HB, JR.HB, or ERET instructions. Software must not assume that a fixed latency will clear the execution hazard.

31	26	25	16	15	6	5	0
POOL32A 000000	0 0000000000	ERET 1111001101	POOL32AXf 111100				
6	10	10	6				

Format: ERET

microMIPS

Purpose: Exception Return

To return from interrupt, exception, or error trap.

Description:

ERET clears execution and instruction hazards, conditionally restores $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ in a Release 2 implementation, and returns to the interrupted instruction at the completion of interrupt, exception, or error processing. ERET does not execute the next instruction (that is, it has no delay slot).

Restrictions:

Pre-Release 6: The operation of the processor is **UNDEFINED** if an ERET is executed in the delay slot of a branch or jump instruction. In Release 6, in the absence of delay/forbidden slots, this restriction does not apply.

An ERET placed between an LL and SC instruction will always cause the SC to fail.

ERET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the ERET returns.

In a Release 2 implementation, ERET does not restore $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ if $Status_{BEV} = 1$, or if $Status_{ERL} = 1$ because any exception that sets $Status_{ERL}$ to 1 (Reset, Soft Reset, NMI, or cache error) does not save $SRSCtl_{CSS}$ in $SRSCtl_{PSS}$. If software sets $Status_{ERL}$ to 1, it must be aware of the operation of an ERET that may be subsequently executed.

Operation:

```

if StatusERL = 1 then
    temp ← ErrorEPC
    StatusERL ← 0
else
    temp ← EPC
    StatusEXL ← 0
    if (ArchitectureRevision() ≥ 2) and (SRSCtlHSS > 0) and (StatusBEV = 0) then
        SRSCtlCSS ← SRSCtlPSS
    endif
endif
if IsMIPS16Implemented() | (Config3ISA > 0) then
    PC ← temp63..1 || 0
    ISAMode ← temp0
else
    PC ← temp
endif
LLbit ← 0
ClearHazards()

```

Exceptions:

Coprocessor Unusable Exception

31	26	25	16	15	6	5	0
POOL32A 000000	0 000000000	1	ERET 1111001101		POOL32AXf 111100		
6	9	1	10		6		

Format: ERETNC

microMIPS Release 5

Purpose: Exception Return No Clear

To return from interrupt, exception, or error trap without clearing the LLbit.

Description:

ERETNC clears execution and instruction hazards, conditionally restores $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ when implemented, and returns to the interrupted instruction at the completion of interrupt, exception, or error processing. ERETNC does not execute the next instruction (i.e., it has no delay slot).

ERETNC is identical to ERET except that an ERETNC will not clear the LLbit that is set by execution of an LL instruction, and thus when placed between an LL and SC sequence, will never cause the SC to fail.

An ERET must continue to be used by default in interrupt and exception processing handlers. The handler may have accessed a synchronizable block of memory common to code that is atomically accessing the memory, and where the code caused the exception or was interrupted. Similarly, a process context-swap must also continue to use an ERET in order to avoid a possible false success on execution of SC in the restored context.

Multiprocessor systems with non-coherent cores (i.e., without hardware coherence snooping) should also continue to use ERET, because it is the responsibility of software to maintain data coherence in the system.

An ERETNC is useful in cases where interrupt/exception handlers and kernel code involved in a process context-swap can guarantee no interference in accessing synchronizable memory across different contexts. ERETNC can also be used in an OS-level debugger to single-step through code for debug purposes, avoiding the false clearing of the LLbit and thus failure of an LL and SC sequence in single-stepped code.

Software can detect the presence of ERETNC by reading $Config5_{LLB}$.

Restrictions:

ERETNC implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes. (For Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream.) The effects of this barrier are seen starting with the instruction fetch and decode of the instruction in the PC to which the ERETNC returns.

Operation:

```

if StatusERL = 1 then
    temp ← ErrorEPC
    StatusERL ← 0
else
    temp ← EPC
    StatusEXL ← 0
    if (ArchitectureRevision() ≥ 2) and (SRSCtlHSS > 0) and (StatusBEV = 0) then
        SRSCtlCSS ← SRSCtlPSS
    endif
endif
if IsMIPS16Implemented() | (Config3ISA > 0) then
    PC ← temp63..1 || 0
    ISAMode ← temp0
else

```

```
    PC ← temp
endif
ClearHazards()
```

Exceptions:

Coprocessor Unusable Exception

31	26	25	21	20	16	15	6	5	0
POOL32A 000000	rt		0 00000	EVP 0011100101			POOL32AXf 111100		
6	5		5	10			6		

Format: EVP rt

microMIPS Release 6

Purpose: Enable Virtual Processor

To enable all virtual processors in a physical core other than the virtual processor that issued the instruction.

Description: $GPR[rt] \leftarrow VPControl$; $VPControl_{DIS} \leftarrow 0$

Enabling a virtual processor means that instruction fetch is resumed.

After all other virtual processors have been enabled, $VPControl_{DIS}$ is cleared. Prior to modification, if rt is non-zero, sign-extended $VPControl$ is written to $GPR[rt]$. If EVP is specified without rt , then rt must be 0.

See the DVP instruction to understand the application of EVP in the context of WAIT/PAUSE/external-control (“DVP” on page 246).

The execution of a DVP must be followed by the execution of an EVP. The execution of an EVP causes execution to resume immediately, *where applicable*, on all other virtual processors, as if the DVP had not been executed, that is, execution is completely restorable after the EVP. On the other hand, if an event occurs in between the DVP and EVP that renders state of the virtual processor UNPREDICTABLE (such as power-gating), then the effect of EVP is UNPREDICTABLE.

EVP may only take effect if $VPControl_{DIS}=1$. Otherwise it is treated as a NOP

This instruction never executes speculatively. It must be the oldest unretired instruction to take effect.

This instruction is only available in Release 6 implementations. For implementations that do not support multi-threading ($Config5_{VP}=0$), this instruction must be treated as a NOP instruction.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 6 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

The pseudo-code below assumes that the EVP is executed by virtual processor 0, while the target virtual processor is numbered ‘n’, where n is each of all remaining virtual processors.

```

if (VPControlDIS = 1)

    // Pseudo-code in italics provides recommended action wrt other VPs
    enable_fetch(VPn) {
        if PAUSE(VPn) retires prior or at disable event
            then VPn execution is not resumed if LLbit is cleared prior to EVP
    }
    enable_interrupt(VPn) {
        if WAIT(VPn) retires prior or at disable event
            then interrupts are ignored by VPn until EVP
    }
endif

```

```

data ← VPControl
GPR[rt] ← sign_extend(data)
VPControlDIS ← 0

```

Exceptions:

Coprocessor Unusable

Reserved Instruction (pre-Release 6 implementations)

Programming Notes:

Before executing an EVP in a DVP/EVP pair, software should first read `VPControlDIS`, returned by DVP, to determine whether the virtual processors are already disabled. If so, the DVP/EVP sequence should be abandoned. This step allows software to safely nest DVP/EVP pairs.

Privileged software may use DVP/EVP to disable virtual processors on a core, such as for the purpose of doing a cache flush without interference from other processes in a system with multiple virtual processors or physical cores.


DVP (and EVP) may be used in other cases such as for power-savings or changing state that is applicable to all virtual processors in a core, such as virtual processor scheduling priority, as described below:

```

ll t0 0(a0)
dvp    // disable all other virtual processors
pause  // wait for LLbit to clear
evp    // enable all othe virtual processors

ll t0 0(a0)
dvp    // disable all other virtual processors
<change core-wide state>
evp    // enable all othe virtual processors

```



Format: EXT *rt*, *rs*, *pos*, *size*

microMIPS

Purpose: Extract Bit Field

To extract a bit field from GPR *rs* and store it right-justified into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{ExtractField}(GPR[rs], msbd, lsb)$

The bit field starting at bit *pos* and extending for *size* bits is extracted from GPR *rs* and stored zero-extended and right-justified in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbd* (the most significant bit of the destination field in GPR *rt*), in instruction bits **15..11**, and *lsb* (least significant bit of the source field in GPR *rs*), in instruction bits **10..6**, as follows:


$msbd \leftarrow size - 1$
 $lsb \leftarrow pos$

The values of *pos* and *size* must satisfy all of the following relations:

$0 \leq pos < 32$
 $0 < size \leq 32$
 $0 < pos + size \leq 32$

Figure 3-9 shows the symbolic operation of the instruction.

Figure 5.11 Operation of the EXT Instruction



Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

The operation is **UNPREDICTABLE** if $lsb + msbd > 31$.

If GPR *rs* does not contain a sign-extended 32-bit value (bits **63..31** equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

if $((lsb + msbd) > 31)$ or $(\text{NotWordValue}(GPR[rs]))$ then
UNPREDICTABLE


```
endif  
temp ← sign_extend( $0^{32-(msbd+1)}$  || GPR[rs]msbd+1sb..1sb)  
GPR[rt] ← temp
```

Exceptions:

Reserved Instruction

31	26	25	22	21	20	16	15	14	13	6	5	0	
POOL32F 010101			ft		fs		0	fmt	FLOOR.L 00001100			POOL32FXf 111011	
6			5		5		1	1	8			6	

Format: FLOOR.L.fmt
 FLOOR.L.S ft, fs
 FLOOR.L.D ft, fs

MIPS64, microMIPS
MIPS64, microMIPS

Purpose: Floating Point Floor Convert to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding down

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and rounded toward \geq (rounding mode 3). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation Enable bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *ft* for long fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model; it is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


$\text{StoreFPR}(ft, L, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact



Format: FLOOR.W.fmt
 FLOOR.W.S ft, fs
 FLOOR.W.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Floor Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding down

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and rounded toward \rightarrow (rounding mode 3). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{31}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{31}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{31}-1$
- $-2^{31}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{31}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *ft* for word fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:

$\text{StoreFPR}(ft, W, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, W))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact

31	26	25	21	20	16	15	13	12	11	10	6	5	0
POOL32A 000000	00000	rs	011	GINVI 00	00101	POOL32AXf 111100							
6	5	5	3	2	5	6							

Format: GINVI rs

microMIPS Release 6

Purpose: Global Invalidate Instruction Cache

In a multi-processor system, fully invalidate all remote primary instruction-caches, or a specified single cache. The local primary instruction cache is also fully invalidated in the case where all the remote caches are to be invalidated.

Description: Invalidate_All_Primary_Instruction_Caches(null or rs)

Fully invalidate all remote primary instruction caches, or a specified single cache, whether local or remote. The local primary instruction cache is also fully invalidated in the case where all remote caches are to be invalidated.

If rs field of the opcode is 0, then all caches are to be invalidated. ‘rs’ should be specified as 0 in the assembly syntax for this case. If rs field of the opcode is not 0, then a single cache that is specified by an implementation dependent number of lower bits of GPR[rs] is invalidated, which may be the local cache itself.

Software based invalidation of the primary instruction cache is required in a system if coherency of the cache is not maintained in hardware. While typically limited to the primary cache, the scope of the invalidation within a processor is however implementation dependent - it should apply to all instruction caches within the cache hierarchy that required software coherence maintenance.

In legacy systems, it is software’s responsibility to keep the instruction cache state consistent through SYNCI instructions. This instruction provides a method for bulk invalidating the instruction caches in lieu of SYNCI.

The instruction’s action is considered complete when the both the local and remote cache invalidations are complete, that is, the data in the cache is no longer available to the related instruction stream. Whether these invalidations are complete can only be determined by the completion of a SYNC (stype=0x14) that follows the invalidate instruction(s). With the completion of the SYNC operation, all global invalidations preceding the SYNC in the program are considered globally visible.

Whether the SYNC(stype=0x14) or the global invalidate itself cause synchronization of the instruction stream to new state/context is implementation dependent.

A processor may send a global invalidate instruction remotely only when any preceding global invalidate for the program has reached a global ordering point.

The GINVI has no instruction or execution hazard barrier semantics in itself.

If the implementation allows a cache line to be locked, i.e., not replaceable during a fill, GINVI will not invalidate the line. A cache line can be locked through the optional CACHE “Fetch and Lock” instruction.

See **Programming Notes** for programming constraints.

Restrictions:

If an implementation does not support the instruction, a Reserved Instruction exception is caused.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In a single processor SOC, this instruction acts on the local instruction cache only.

Operation:

```
Local:
    if (Config5GI ≠ 2'b1x)
```

```

        SignalException(ReservedInstruction) // if not implemented
        break
    endif
    if IsCoproprocessorEnabled(0) then
        // Fully invalidate local instruction cache, if selected.
        // Send invalidation message to other cores, if required.
    else
        SignalException(CoproprocessorUnusable, 0)
    endif

    Remote:
        // Fully invalidate remote instruction cache.

```

Exceptions:

Reserved Instruction, Coprocessor Unusable

Programming Notes:

For the local processor, the instruction stream is synchronized by an instruction hazard barrier such as JR.HB.

The instruction stream in the remote processor is synchronized with respect to the execution of GINVI once the SYNC operation following GINVI completes.

The following sequence is recommended for use of GINVI.

```

ginvi      /* fully-invalidate all caches*/
sync 0x14 /* Enforce completion - all instruction streams synchronized. */
jr.hb ra  /* Clear instruction hazards*/

```

Implementation Notes:

None.

31	26	25	21	20	16	15	13	12	11	10	9	8	6	5	0
POOL32A 000000	00000	rs	011	GINVT 10	type	101	POOL32AXf 111100								
6	5	5	3	2	2	3	6								

Format: GINVT rs, type

microMIPS Release 6

Purpose: Global Invalidate TLB

In a multi-processor system, invalidate translations of remote TLBs and local TLB.

Description: `Invalidate_TLB(GPR[rs], MemoryMapID)`

Invalidate a TLB in multiple ways - entire TLB, by Virtual Address and MemoryMapID, by MemoryMapID or Virtual Address. The Virtual Address is obtained from GPR[rs]. The MemoryMapID is derived from CP0 *MemoryMapID*. The virtual address is associated with a specific Memory Map identified by MemoryMapID.

The virtual address within GPR[rs] is aligned to twice the size of the minimum page size of 4KB i.e., it is equivalent to *EntryHi_VPN2*: bit 13 of the virtual address is aligned to bit 13 of GPR[rs]. If the virtual address is not required, such as in the case of invalidate All or by MemoryMapID, then 'rs' should be specified as 0 in the assembly syntax but is otherwise ignored.

The MemoryMapID is a replacement for *EntryHi_ASID*. The MemoryMapID is an implementation-dependent number of bits that must be larger than the existing *EntryHi_ASID* (10-bits including *EntryHi_ASIDX*). The purpose of a larger tag is to be able to uniquely identify processes in the system. A 16-bit MemoryMapID for example will identify 64K Memory Maps, while the current 8-bit ASID only identifies 256, and is thus subject to frequent recycling. An implementation with MemoryMapID is designed to be backward compatible with software that uses *EntryHi_ASID*. See CP0 *MemoryMapID*.

Table 5.1 specifies the different types of invalidates supported as a function of the "type" field of the instruction.

Table 5.1 Types of Global TLB Invalidates

Encoding of "type" field	Definition
00	Invalidate entire TLB
01	Invalidate by VA (MemoryMapID is globalized)
10	Invalidate by MemoryMapID
11	Invalidate by VA and MemoryMapID.

With reference to Table 5.1, if the Global bit in a TLB entry is set, then MemoryMapID comparison is ignored by the operation.

The instruction is considered complete when the local and remote invalidations are complete. Whether these invalidations are complete can only be determined by the completion of a SYNC (stype=0x14) that follows the invalidate instruction(s). With the completion of the SYNC operation, all invalidations of this type preceding the SYNC in the program are considered globally visible.

Whether the SYNC(stype=0x14) or the global invalidate itself cause synchronization of the instruction stream to new state/context is implementation dependent.

A GINVT based invalidation is complete, whether local or remote, when the following has occurred: the TLB is invalidated of matching entries, and all instructions in the instruction stream after the point of completion can only access the new context.

A processor may send a global invalidate instruction remotely only when any preceding global invalidate for the program has reached a global ordering point.

GINVT has no instruction or execution hazard barrier semantics in itself.

A GINVT operation that is specified to invalidate all entries will only invalidate non-wired entries. Other GINVT operations will invalidate wired entries on a match.

Restrictions:

If an implementation does not support the instruction, or use of MemoryMapID is disabled, then a Reserved Instruction exception is caused.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable exception is signaled.

In a single processor SOC, this instruction acts on the local TLB only.

Operation:

```

Local:
  if (Config5GI ≠ 2'b11) then
    SignalException(ReservedInstruction, 0)
    break
  endif
  if IsCoprocessorEnabled(0) then
    if (Config5MI = 1)
      // generate control from instruction encoding.
      invAll ← (ginvt[type] = 0b00)
      invVA ← (ginvt[type] = 0b01)
      invMMid ← (ginvt[type] = 0b10)
      invVAMMid ← (ginvt[type] = 0b11)
      // generate data; how data is driven when unsupported is imp-dep.
      // Format of GPR[rs] equals CP0 EntryHi.
      InvMsgVPN2 ← GPR[rs]VPN2msb..13           // VPN2msb is imp-dep in MIPS64
      InvMsgR ← GPR[rs]63:62                   // R same as CP0 EntryLo.R
      InvMsgMMid ← MemoryMapID                  // imp-dep # of bits
      // Broadcast invalidation message to other cores.
      InvalidateTLB(InvMsgVPN2, InvMsgR, InvMsgMMid, invAll, invVAMMid, invMMid, invVA)
    else // if not implemented, MMid disabled
      SignalException(ReservedInstruction)
    endif
  else
    SignalException(CoprocessorUnusable, 0)
  endif

Remote:
  // Repeat in all remote TLBs
  InvalidateTLB(InvMsgVPN2, InvMsgR, InvMsgMMid, invAll, invVAMMid, invMMid, invVA)

function InvalidateTLB(InvMsgVPN2, InvMsgR, InvMsgMMid, invAll, invVAMMid, invMMid, invVA)
  // "Mask" is equivalent to CP0 PageMask.
  // "G" is equivalent to the Global bit in CP0 EntryLo0/1.
  // "R" is equivalent to the R bit in CP0 EntryHi.
  for i in 0..TLBEntries-1
    // Wired entries are excluded.
    VAMatch ← ((TLB[i]VPN2 and not TLB[i]Mask) = (InvMsgVPN2 and not
TLB[i]Mask)) and (TLB[i]R = InvMsgR)
    MMidMatch ← (TLB[i]MMid = InvMsgMMid)
    if ((invAll and (i > CP0.Wired.Wired)) or // do not invalidate Wired

```

```

        (VAMatch and ((TLB[i]G = 1) or MMidMatch) and invVAMMid) or
        (VAMatch and invVA) or
        (MMidMatch and (TLB[i]G ≠ 1) and invMMid)) then
            TLB[i]HW_Valid ← 0      // where HW_Valid is the entry valid bit
        endif
    endfor
endfunction

```

Exceptions:

Reserved Instruction, Coprocessor Unusable

Programming Notes:

Since CP0 *MemoryMapID* sources the value of MemoryMapID of the currently running process, the kernel must save/restore *MemoryMapID* appropriately before it modifies it for the invalidation operation. Between the save and restore, it must utilize unmapped addresses.

An MTC0 that modifies *MemoryMapID* must be followed by an EHB to make this value visible to a subsequent GINVT. Where multiple GINVTs are used prior to a single SYNC (stype=0x14), each may use a different value of *MemoryMapID*.

For the local processor, the instruction stream is synchronized to the new translation context (where applicable) by an instruction hazard barrier such as JR.HB.

The instruction stream in the remote processor is synchronized with respect to the execution of GINVT once the SYNC operation completes.

The following sequence is recommended for use of GINVT.

```


mtc0 0, CO_PWctl /* disable Page Walker, where applicable; implementation-dependent */
ehb              /* Clear execution hazards to prevent speculative walks */
ginv t r1, type  /* Invalidate TLB(s) */
sync 0x14        /* Enforce completion */
jr.hb ra         /* Clear instruction hazards */

```

Whether the hardware page table walker, if implemented, needs to be disabled as shown above, is implementation dependent. It is recommended that hardware take the steps to locally disable the hardware page table walker to maintain TLB consistency, as it would for remote TLBs.

Software must take into account a system that may have potentially varying widths of *MemoryMapID*. While not recommended, different processors may have different implemented or programmed widths. Further, the interface between processors may support yet another width. If this is the case, then software responsible for global invalidates should be run on the processor with maximum width. Software must zero-fill any bits that are unused by a target. Software should also be able to rely on the implementation zero-filling bits where widths increase across any interface.

If an intermediate interface between source and target truncates the width of *MemoryMapID*, then software could address this limitation through various means: It could restrict the use of *MemoryMapID* to the interface width, it could program *MemoryMapID* with the expectation that over-invalidation may occur, or it should default to legacy means of invalidating the caches to prevent unreliable system behavior.



Format: INS *rt*, *rs*, *pos*, *size*

microMIPS

Purpose: Insert Bit Field

To merge a right-justified bit field from GPR *rs* into a specified field in GPR *rt*.

Description: $GPR[rt] \leftarrow InsertField(GPR[rt], GPR[rs], msb, lsb)$

The right-most *size* bits from GPR *rs* are merged into the value from GPR *rt* starting at bit position *pos*. The result is placed back in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msb* (the most significant bit of the field), in instruction bits 15..11, and *lsb* (least significant bit of the field), in instruction bits 10..6, as follows:


$msb \leftarrow pos + size - 1$
 $lsb \leftarrow pos$

The values of *pos* and *size* must satisfy all of the following relations:

$0 \leq pos < 32$
 $0 < size \leq 32$
 $0 < pos + size \leq 32$

Figure 3-10 shows the symbolic operation of the instruction.

Figure 5.12 Operation of the INS Instruction



Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

The operation is **UNPREDICTABLE** if $lsb > msb$.

If either GPR rs or GPR rt does not contain sign-extended 32-bit values (bits **63..31** equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

    if (lsb > msb) or (NotWordValue(GPR[rs])) or (NotWordValue(GPR[rt])) then
        UNPREDICTABLE
    endif
    GPR[rt] ← sign_extend(GPR[rt]31..msb+1 || GPR[rs]msb-lsb..0 || GPR[rt]lsb-1..0)

```

Exceptions:

Reserved Instruction



Format: JALRC rs (rt = 31 implied)
JALRC rt, rs

microMIPS Release 6
microMIPS Release 6

Purpose: Jump and Link Register Compact

To execute a procedure call to an instruction address in a register

Description: $GPR[rt] \leftarrow return_addr$, $PC \leftarrow GPR[rs]$

Place the return address link in GPR *rt*. The return link is the address of the first instruction following the branch, where execution continues after a procedure call.

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

Restrictions Related to Multiple Instruction Sets: This instruction can change the active instruction set, if more than one instruction set is implemented.

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors which implement MIPS and if the ISAMode bit of the target is MIPS (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS ISA, if the intended target ISAMode is MIPS (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump; that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

Release 6 maps JR and JR.HB to JALRC and JALRC.HB with *rt* = 0:

Release 6 assemblers should accept the JR and JR.HB mnemonics, mapping them to the Release 6 instruction encodings.

Operation:

$temp \leftarrow GPR[rs]$


```
GPR[rt] ← PC + 4
if (Config3ISA = 1) then
    PC ← temp
else
    PC ← tempGPREN-1..1 || 0
    ISAMode ← temp0
endif
```

Exceptions:

None

Programming Notes:

This jump-and-link register instruction can select a register for the return link; other link instructions use GPR 31. The default register for GPR *rd*, if omitted in the assembly language instruction, is GPR 31.



Format: JALRC.HB *rs* (*rt* = 31 implied)
JALRC.HB *rt*, *rs*

microMIPS Release 6
microMIPS Release 6

Purpose: Jump and Link Register Compact with Hazard Barrier

To execute a procedure call to an instruction address in a register and clear all execution and instruction hazards

Description: $GPR[rt] \leftarrow return_addr$, $PC \leftarrow GPR[rs]$, clear execution and instruction hazards

Place the return address link in GPR *rt*. The return link is the address of the second instruction following the branch, where execution continues after a procedure call.

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

JALRC.HB implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the JALRC.HB instruction jumps. An equivalent barrier is also implemented by the ERET instruction, but that instruction is only available if access to Coprocessor 0 is enabled, whereas JALRC.HB is legal in all operating modes.

This instruction clears both execution and instruction hazards. Refer to the [EHB](#) instruction description for the method of clearing execution hazards alone.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

After modifying an instruction stream mapping or writing to the instruction stream, execution of those instructions has **UNPREDICTABLE** behavior until the instruction hazard has been cleared with JALRC.HB, JALRSC.HB, JR.HB, ERET, or DERET. Further, the operation is **UNPREDICTABLE** if the mapping of the current instruction stream is modified.

Restrictions Related to Multiple Instruction Sets: This instruction can change the active instruction set, if more than one instruction set is implemented.

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors which implement MIPS and if the ISAMode bit of the target address is MIPS (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS ISA, if the intended target ISAMode is MIPS (bit 0 of GPR *rs* is zero), an

Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump; that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

Release 6 maps JR and JR.HB to JALRC and JALRC.HB with $rt = 0$:

Release 6 assemblers should accept the JR and JR.HB mnemonics, mapping them to the Release 6 instruction encodings.

Operation:

```
temp ← GPR[rs]
GPR[rt] ← PC + 4
if (Config3ISA = 1) then
    PC ← temp
else
    PC ← tempGPRLEN-1..1 || 0
    ISAMode ← temp0
endif
ClearHazards()
```

Exceptions:

None

Programming Notes:


This branch-and-link instruction can select a register for the return link; other link instructions use GPR 31. The default register for GPR rt , if omitted in the assembly language instruction, is GPR 31.

Release 6 JR.HB rs is implemented as JALRC.HB $r0, rs$. For example, as JALRC.HB with the destination set to the zero register, $r0$.

This instruction implements the final step in clearing execution and instruction hazards before execution continues. A hazard is created when a Coprocessor 0 or TLB write affects execution or the mapping of the instruction stream, or after a write to the instruction stream. When such a situation exists, software must explicitly indicate to hardware that the hazard should be cleared. Execution hazards alone can be cleared with the EHB instruction. Instruction hazards can only be cleared with a JR.HB, JALRC.HB, or ERET instruction. These instructions cause hardware to clear the hazard before the instruction at the target of the jump is fetched. Note that because these instructions are encoded as jumps, the process of clearing an instruction hazard can often be included as part of a call (JALR) or return (JR) sequence, by simply replacing the original instructions with the HB equivalent.

Example: Clearing hazards due to an ASID change

```
/*
 * Code used to modify ASID and call a routine with the new
 * mapping established.
 */
 * a0 = New ASID to establish
 * a1 = Address of the routine to call
 */
mfc0    v0, C0_EntryHi      /* Read current ASID */
li      v1, ~M_EntryHiASID /* Get negative mask for field */
and     v0, v0, v1          /* Clear out current ASID value */
or      v0, v0, a0          /* OR in new ASID value */
mtc0    v0, C0_EntryHi      /* Rewrite EntryHi with new ASID */
jalrc.hb a1                 /* Call routine, clearing the hazard */
```

Format: JIALC rt, offset

microMIPS32 Release 6

Purpose: Jump Indexed and Link, Compact

Description: $GPR[31] \leftarrow PC+4$, $PC \leftarrow (GPR[rt] + \text{sign_extend}(\text{offset}))$

The jump target is formed by sign extending the offset field of the instruction and adding it to the contents of GPR rt.

The offset is NOT shifted, that is, each bit of the offset is added to the corresponding bit of the GPR.

Places the return address link in GPR 31. The return link is the address of the following instruction, where execution continues after a procedure call returns. Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

Compact jumps do not have delay slots. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors that implement MIPS and if the ISAMode bit of the target is MIPS (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS ISA, if the intended target ISAMode is MIPS (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Exceptions:

None

Operation:


```
temp ← GPR[rt] + sign_extend(offset)
GPR[31] ← PC + 4
if (Config3_ISA = 1) then
```



```
    PC ← temp
else
    PC ← (tempGPRLEN-1..1 || 0)
    ISAMode ← temp0
endif
```

Programming Notes:

JIALC does NOT shift the offset before adding it the register. This can be used to eliminate tags in the least significant bits that would otherwise produce misalignment. It also allows JIALC to be used as a substitute for the JALX instruction, removed in Release 6, where the lower bits of the target PC, formed by the addition of GPR[*rt*] and the unshifted offset, specify the target ISAMode.



Format: JIC rt, offset

microMIPS32 Release 6

Purpose: Jump Indexed, Compact

Description: $PC \leftarrow (GPR[rt] + \text{sign_extend}(\text{offset}))$

The branch target is formed by sign extending the offset field of the instruction and adding it to the contents of GPR rt.

The offset is NOT shifted, that is, each bit of the offset is added to the corresponding bit of the GPR.

For processors that do not implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Bit 0 of GPR *rs* is interpreted as the target ISA Mode: if this bit is 0, signal an Address Error exception when the target instruction is fetched because this target ISA Mode is not supported. Otherwise, set bit 0 of the target address to zero, and fetch the instruction.

For processors that do implement the MIPS64 ISA:

- Jump to the effective target address in GPR *rs*. Set the ISA Mode bit to the value in GPR *rs* bit 0. Set bit 0 of the target address to zero. If the target ISA Mode bit is 0 and the target address is not 4-byte aligned, an Address Error exception will occur when the target instruction is fetched.

Compact jumps do not have a delay slot. The instruction after the jump is NOT executed when the jump is executed.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 of GPR *rs*.

For processors that implement MIPS64 and if the ISAMode bit of the target is MIPS64 (bit 0 of GPR *rs* is 0) and address bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

For processors that do not implement MIPS64 ISA, if the intended target ISAMode is MIPS64 (bit 0 of GPR *rs* is zero), an Address Error exception occurs when the jump target is fetched as an instruction.

Any instruction, including a branch or jump, may immediately follow a branch or jump, that is, delay slot restrictions do not apply in Release 6.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Exceptions:

None


Operation:

```
temp ← GPR[rt] + sign_extend(offset)
if (Config3ISA = 1) then
    PC ← temp
else
    PC ← (tempGPREN-1..1 || 0)
```

```
    ISAMode ← temp0  
endif
```

Programming Notes:

JIC does NOT shift the offset before adding it the register. This can be used to eliminate tags in the least significant bits that would otherwise produce misalignment. It also allows JIALC to be used as a substitute for the JALX instruction, removed in Release 6, where the lower bits of the target PC, formed by the addition of GPR[rt] and the unshifted offset, specify the target ISAMode.



Format: LB *rt*, *offset* (*base*)

microMIPS

Purpose: Load Byte

To load a byte from memory as a signed value.

Description: $GPR[rt] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, sign-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
memdoubleword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor BigEndianCPU3
GPR[rt] ← sign_extend(memdoubleword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LBE rt, offset(base)

microMIPS

Purpose: Load Byte EVA

To load a byte as a signed value from user mode virtual address space when executing in kernel mode.

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, sign-extended, and placed in GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The LBE instruction functions the same as the LB instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode and executing in kernel mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

Only usable when access to Coprocessor0 is enabled and accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
memdoubleword ← LoadMemory(CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor BigEndianCPU3
GPR[rt] ← sign_extend(memdoubleword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid

Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



Format: LBU *rt*, *offset*(*base*)

microMIPS

Purpose: Load Byte Unsigned

To load a byte from memory as an unsigned value

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
memdoubleword ← LoadMemory(CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor BigEndianCPU3
GPR[rt] ← zero_extend(memdoubleword_7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LBUE rt, offset(base)

microMIPS

Purpose: Load Byte Unsigned EVA

To load a byte as an unsigned value from user mode virtual address space when executing in kernel mode.

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, zero-extended, and placed in GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The LBUE instruction functions the same as the LBU instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

Only usable when access to Coprocessor0 is enabled and accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor ReverseEndian3)
memdoubleword ← LoadMemory(CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor BigEndianCPU3
GPR[rt] ← zero_extend(memdoubleword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



Format: LD *rt*, *offset*(*base*)

microMIPS64

Purpose: Load Doubleword

To load a doubleword from memory

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
memdoubleword ← LoadMemory (CCA, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdoubleword

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Reserved Instruction, Watch



Format: LDC1 ft, offset(base)

microMIPS

Purpose: Load Doubleword to Floating Point

To load a doubleword from memory to an FPR.

Description: $\text{FPR}[\text{ft}] \leftarrow \text{memory}[\text{GPR}[\text{base}] + \text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in FPR *ft*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
memdoubleword ← LoadMemory(CCA, DOUBLEWORD, pAddr, vAddr, DATA)
StoreFPR(ft, UNINTERPRETED_DOUBLEWORD, memdoubleword)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, Address Error, Watch



Format: LDC2 rt, offset(base)

microMIPS

Purpose: Load Doubleword to Coprocessor 2

To load a doubleword from memory to a Coprocessor 2 register.

Description: $CPR[2,rt,0] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in Coprocessor 2 register *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
memdoubleword ← LoadMemory(CCA, DOUBLEWORD, pAddr, vAddr, DATA)
CPR[2,rt,0] ← memdoubleword

```


Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, Address Error, Watch

Programming Notes:

There is no Load Linked Word Unsigned operation corresponding to Load Word Unsigned.

Release 6 implements a 9-bit offset, whereas all release levels lower than Release 6 implement a 16-bit offset.



Format: LDPC rt, offset

microMIPS64 Release 6

Purpose: Load Doubleword PC-relative

To load a doubleword from memory, using a PC-relative address.

Description: $GPR[rt] \leftarrow memory[(PC \& \sim 0x7) + sign_extend(offset \ll 3)]$

The bit offset is shifted left by 3 bits, sign-extended, and added to the address of the aligned doubleword containing the LDPC instruction.

The contents of the 64-bit doubleword at the memory location specified by the effective address are fetched, and placed in GPR *rt*.

Restrictions:

LDPC is naturally aligned, by specification.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Operation

```
vAddr ← ( (PC&~0x7)+ sign_extend(offset) )
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
memdoubleword ← LoadMemory (CCA, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdoubleword
```


Exceptions:

TLB Refill, TLB Invalid, TLB Read Inhibit, Bus Error, Address Error, Watch, Reserved Instruction

Programming Note

The Release 6 PC-relative loads (LWPC, LWUPC, LDPC) are considered data references.

For the purposes of watchpoints (provided by the CP0 *WatchHi* and *WatchLo* registers) and EJTAG breakpoints, the PC-relative reference is considered to be a data reference, rather than an instruction reference. That is, the watchpoint or breakpoint is triggered only if enabled for data references.



Format: LH *rt*, *offset*(*base*)

microMIPS

Purpose: Load Halfword

To load a halfword from memory as a signed value

Description: $GPR[rt] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, sign-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian2 || 0))
memdoubleword ← LoadMemory(CCA, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← sign_extend(memdoubleword15+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LHE *rt*, *offset*(*base*)

microMIPS

Purpose: Load Halfword EVA

To load a halfword as a signed value from user mode virtual address space when executing in kernel mode.

Description: $GPR[rt] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, sign-extended, and placed in GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The LHE instruction functions the same as the LH instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

Only usable when access to Coprocessor0 is enabled and accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian2 || 0))
memdoubleword ← LoadMemory (CCA, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← sign_extend(memdoubleword15+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error

Watch, Reserved Instruction, Coprocessor Unusable



Format: LHU *rt*, *offset*(*base*)

microMIPS

Purpose: Load Halfword Unsigned

To load a halfword from memory as an unsigned value

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian2 || 0))
memdoubleword ← LoadMemory (CCA, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← zero_extend(memdoubleword15+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LHUE rt, offset(base)

microMIPS

Purpose: Load Halfword Unsigned EVA

To load a halfword as an unsigned value from user mode virtual address space when executing in kernel mode.

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The LHUE instruction functions the same as the LHU instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

Only usable when access to Coprocessor0 is enabled and accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian2 || 0))
memdoubleword ← LoadMemory(CCA, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← zero_extend(memdoubleword15+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



Format: LL *rt*, *offset* (*base*)

microMIPS

Purpose: Load Linked Word

To load a word from memory for an atomic read-modify-write

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The LL and SC instructions provide the primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length, and written into GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

This begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LL is executed it starts an active RMW sequence replacing any other sequence that was active. The RMW sequence is completed by a subsequent SC instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Executing LL on one processor does not cause an action that, by itself, causes an SC for the same block to fail on another processor.

An execution of LL does not have to be followed by execution of SC; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the SC instruction for the formal definition.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Availability and Compatibility:

This instruction has been reallocated an opcode in Release 6.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)

```



```
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)  
LLbit ← 1
```


Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch

Programming Notes:

There is no Load Linked Word Unsigned operation corresponding to Load Word Unsigned.

Release 6 implements a 9-bit offset, whereas all release levels lower than Release 6 implement a 16-bit offset.



Format: LLD *rt*, *offset*(*base*)

microMIPS64

Purpose: Load Linked Doubleword

To load a doubleword from memory for an atomic read-modify-write

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The LLD and SCD instructions provide primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed into GPR *rt*. The signed *offset* is added to the contents of GPR *base* to form an effective address.

This begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LLD is executed it starts the active RMW sequence and replaces any other sequence that was active. The RMW sequence is completed by a subsequent SCD instruction that either completes the RMW sequence atomically and succeeds, or does not complete and fails.

Executing LLD on one processor does not cause an action that, by itself, would cause an SCD for the same block to fail on another processor.

An execution of LLD does not have to be followed by execution of SCD; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the [SCD](#) instruction for the formal definition.

The effective address must be naturally-aligned. If any of the 3 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
memdoubleword ← LoadMemory(CCA, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdoubleword
LLbit ← 1

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Watch



Format: LLE rt, offset(base)

microMIPS

Purpose: Load Linked Word EVA

To load a word from a user mode virtual address when executing in kernel mode for an atomic read-modify-write

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The LLE and SCE instructions provide the primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations using user mode virtual addresses while executing in kernel mode.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length, and written into GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

This begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LLE is executed it starts an active RMW sequence replacing any other sequence that was active. The RMW sequence is completed by a subsequent SCE instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Executing LLE on one processor does not cause an action that, by itself, causes an SCE for the same block to fail on another processor.

An execution of LLE does not have to be followed by execution of SCE; a program is free to abandon the RMW sequence without attempting a write.

The LLE instruction functions the same as the LL instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Segmentation Control for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the **SCE** instruction for the formal definition.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)

```

```
byte ← vAddr2..0 xor (BigEndianCPU || 02)  
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)  
LLbit ← 1
```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Watch, Coprocessor Unusable

Programming Notes:

There is no Load Linked Word Unsigned operation corresponding to Load Word Unsigned.

31	26	25	21	20	16	15	12	11	9	8	4	3	0
POOL32C 011000		rt		base		LLDP 0101		0 000		rd		0 0000	
6		5		5		4		3		5		4	

Format: LLDP *rt*, *rd*, (*base*)

microMIPS64 Release 6

Purpose: Load Linked DoubleWord Paired

To load two double-words from memory for an atomic read-modify-write, writing a double-word each to two registers.

Description: $GPR[rd] \leftarrow memory[GPR[base]]_{127..64}$, $GPR[rt] \leftarrow memory[GPR[base]]_{63..0}$

The LLDP and SCDP instructions provide primitives to implement a paired double-word atomic read-modify-write (RMW) operation at a synchronizable memory location.

The paired double-word at the memory location specified by the quad-word aligned effective address is read in a single atomic memory operation. The least significant double-word is written into GPR *rt*. The most significant double-word is written into GPR *rd*.

A paired double-word read or write occurs as a pair of double-word reads or writes that is quad-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The execution of LLDP begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LLDP is executed it starts an active RMW sequence replacing any other sequence that was active. The RMW sequence is completed by a subsequent SCDP instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Successful execution of the LLDP results in setting LLbit and writing CP0 *LLAddr*, where LLbit is the least-significant bit of *LLAddr*. *LLAddr* contains the data-type aligned address of the operation, in this case a quad-word aligned address.

Executing LLDP on one processor does not cause an action that, by itself, causes a store conditional instruction type for the same block to fail on another processor.

An execution of LLDP does not have to be followed by execution of SCDP; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the **SC** instruction for the formal definition.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if *Config5_{XNP}*=0.

Operation:

$vAddr \leftarrow GPR[base]$

```

(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
// PAIREDDOUBLEWORD: two double-word data-type that is quad-word atomic
memquadword ← LoadMemory (CCA, PAIREDDOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memquadword63..0
GPR[rd] ← memquadword127..64
LLAddr ← pAddr // quad-word aligned i.e., pAddr3..0 are 0, or not supported.
LLbit ← 1


```

Exceptions:

TLB Refill, TLB Invalid, Reserved Instruction, Address Error, Watch

Programming Notes:

An LLDP instruction for which the two destination registers are the same but non-zero is UNPREDICTABLE. An LLDP with two zero destination registers followed by a SCDP can be used to accomplish a quad-word atomic write.



Format: LLWP *rt*, *rd*, (*base*)

microMIPS Release 6

Purpose: Load Linked Word Paired

To load two words from memory for an atomic read-modify-write, writing a word each to two registers.

Description: $GPR[rd] \leftarrow memory[GPR[base]]_{63..32}$, $GPR[rt] \leftarrow memory[GPR[base]]_{31..0}$

The LLWP and SCWP instructions provide primitives to implement a paired word atomic read-modify-write (RMW) operation at a synchronizable memory location.

The 64-bit paired word, as a concatenation of two words, at the memory location specified by the double-word aligned effective address is read. The least significant word, sign-extended to the GPR register length, is written into GPR *rt*, and the most significant word, sign-extended to the GPR register length, is written into GPR *rd*.

A paired word read or write occurs as a pair of word reads or writes that is double-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The execution of LLWP begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LLWP is executed it starts an active RMW sequence replacing any other sequence that was active. The RMW sequence is completed by a subsequent SCWP instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Successful execution of the LLWP results in setting LLbit and writing CP0 *LLAddr*, where LLbit is the least-significant bit of *LLAddr*. *LLAddr* contains the data-type aligned address of the operation, in this case a double-word.

Executing LLWP on one processor does not cause an action that, by itself, causes a store conditional instruction type for the same block to fail on another processor.

An execution of LLWP does not have to be followed by execution of SCWP; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the **SC** instruction for the formal definition.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if *Config5_{XP}*=0.

Operation:

```
vAddr ← GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
```




```
// PAIREDWORD: two word data-type that is double-word atomic
memdoubleword ← LoadMemory (CCA, PAIREDWORD, pAddr, vAddr, DATA)
GPR[rt] ← sign_extend(memdoubleword31..0)
GPR[rd] ← sign_extend(memdoubleword63..32)
LLAddr ← pAddr // double-word aligned i.e., pAddr2..0 are 0, or not supported.
LLbit ← 1
```

Exceptions:

TLB Refill, TLB Invalid, Reserved Instruction, Address Error, Watch

Programming Notes:

An LLWP instruction for which the two destination registers are the same but non-zero is UNPREDICTABLE. An LLWP with two zero destination registers followed by a SCWP can be used to accomplish a double-word atomic write.



Format: LLWPE *rt*, *rd*, (*base*)

microMIPS Release 6

Purpose: Load Linked Word Paired EVA

To load two words from memory for an atomic read-modify-write, writing a word each to two registers. The load occurs in kernel mode from user virtual address space.

Description: $GPR[rd] \leftarrow memory[GPR[base]]_{63..32}$, $GPR[rt] \leftarrow memory[GPR[base]]_{31..0}$

The LLWPE and SCWPE instructions provide primitives to implement a paired word atomic read-modify-write (RMW) operation at a synchronizable memory location.

The 64-bit paired word at the memory location specified by the double-word aligned effective address is read. The least significant word, sign-extended to the GPR register length, is written into GPR *rt*. The most significant word, sign-extended to the GPR register length, is written into GPR *rd*.

A paired word read or write occurs as a pair of word reads or writes that is double-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The execution of LLWPE begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LLWPE is executed it starts an active RMW sequence replacing any other sequence that was active. The RMW sequence is completed by a subsequent SCWPE instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Successful execution of the LLWPE results in setting LLbit and writing CP0 *LLAddr*, where LLbit is the least-significant bit of *LLAddr*. *LLAddr* contains the data-type aligned address of the operation, in this case a double-word aligned address.

The LLWPE instruction functions the same as the LLWP instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Segmentation Control for additional information.

Executing LLWPE on one processor does not cause an action that, by itself, causes a store conditional instruction type for the same block to fail on another processor.

An execution of LLWPE does not have to be followed by execution of SCWPE; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the [SC](#) instruction for the formal definition.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if $Config5_{XNP}=0$ and $Config5_{EVA}=1$.

Operation:

```

vAddr ← GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
// PAIREDWORD: two word data-type that is double-word atomic
memdoubleword ← LoadMemory (CCA, PAIREDWORD, pAddr, vAddr, DATA)
GPR[rt] ← sign_extend(memdoubleword31..0)
GPR[rd] ← sign_extend(memdoubleword63..32)
LLAddr ← pAddr // double-word aligned i.e., pAddr2..0 are 0, or not supported.
LLbit ← 1

```

Exceptions:

TLB Refill, TLB Invalid, Reserved Instruction, Address Error, Watch, Coprocessor Unusable.

Programming Notes:

An LLWPE instruction for which the two destination registers are the same but non-zero is UNPREDICTABLE. An LLWPE with two zero destination registers followed by a SCWPE can be used to accomplish a double-word atomic write.

31	26	25	21	20	16	15	11	10	9	8	6	5	0	
POOL32A 000000	rt		rs		rd			sa		000		LSA 001111		
POOL32S 010110	rt		rs		rd			sa		100		DLSA 000100		
6		5		5		5			2		3		6	

Format: LSA DLSA
 LSA rt, rs, rd, sa
 DLSA rt, rs, rd, sa

microMIPS32 Release 6
microMIPS64 Release 6

Purpose: Load Scaled Address, Doubleword Load Scaled Address

Description:

LSA: $\text{GPR}[\text{rd}] \leftarrow \text{sign_extend.32}(\text{GPR}[\text{rs}] \ll (\text{sa}+1) + \text{GPR}[\text{rt}])$
 DLSA: $\text{GPR}[\text{rd}] \leftarrow (\text{GPR}[\text{rs}] \ll (\text{sa}+1) + \text{GPR}[\text{rt}])$

LSA adds two values derived from registers *rs* and *rt*, with a scaling shift on *rs*. The scaling shift is formed by adding 1 to the 2-bit *sa* field, which is interpreted as unsigned. The scaling left shift varies from 1 to 5, corresponding to multiplicative scaling values of $\times 2$, $\times 4$, $\times 8$, $\times 16$, bytes, or 16, 32, 64, or 128 bits.

LSA is a MIPS32 compatible instruction, sign extending its result from bit 31 to bit 63.

DLSA is a MIPS64 compatible instruction, performing the scaled index calculation fully 64-bits wide.

Restrictions:

LSA: None

DLSA: Reserved Instruction exception if 64-bit instructions are not enabled.

Availability and Compatibility:

LSA instruction is introduced by and required as of Release 6.

DLSA instruction is introduced by and required as of Release 6.


Operation

LSA: $\text{GPR}[\text{rd}] \leftarrow \text{sign_extend.32}(\text{GPR}[\text{rs}] \ll (\text{sa}+1) + \text{GPR}[\text{rt}])$
 DLSA: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rs}] \ll (\text{sa}+1) + \text{GPR}[\text{rt}]$

Exceptions:

LSA: None

DLSA: Reserved Instruction



Format: LUI *rt*, *immediate*

microMIPS, Assembly Idiom Release 6

Purpose: Load Upper Immediate

To load a constant into the upper half of a word

Description: $\text{GPR}[\text{rt}] \leftarrow \text{sign_extend}(\text{immediate} \parallel 0^{16})$

The 16-bit *immediate* is shifted left 16 bits and concatenated with 16 bits of low-order zeros. The 32-bit result is sign-extended and placed into GPR *rt*.

Restrictions:

None.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:


$\text{GPR}[\text{rt}] \leftarrow \text{sign_extend}(\text{immediate} \parallel 0^{16})$

Exceptions:

None

Programming Notes:

In Release 6, LUI is an assembly idiom of AUI with *rs*=0.



Format: LW *rt*, *offset* (*base*)

microMIPS

Purpose: Load Word

To load a word from memory as a signed value

Description: $GPR[rt] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr_2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LWC1 ft, offset(base)

microMIPS

Purpose: Load Word to Floating Point

To load a word from memory to an FPR

Description: $\text{FPR}[ft] \leftarrow \text{memory}[\text{GPR}[base] + \text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and placed into the low word of FPR *ft*. If FPRs are 64 bits wide, bits 63..32 of FPR *ft* become **UNPREDICTABLE**. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
bytesel ← vAddr_2..0 xor (BigEndianCPU || 02)
StoreFPR(ft, UNINTERPRETED_WORD, memdoubleword_31+8*bytesel..8*bytesel)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



Format: LWC2 rt, offset(base)

microMIPS

Purpose: Load Word to Coprocessor 2

To load a word from memory to a COP2 register.

Description: $CPR[2,rt,0] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and placed into the low word of **COP2** (Coprocessor 2) general register *rt*. The signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{+EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory(CCA, DOUBLEWORD, pAddr, vAddr, DATA)
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
CPR[2,rt,0] ← sign_extend(memdoubleword31+8*bytesel..8*bytesel)


```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch

Programming Notes:

Release 6 implements an 11-bit offset, whereas all release levels lower than Release 6 implement a 16-bit offset.



Format: LWE rt, offset(base)

microMIPS

Purpose: Load Word EVA

To load a word from user mode virtual address space when executing in kernel mode.

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The LWE instruction functions the same as the LW instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Restrictions:

Only usable when access to Coprocessor0 is enabled and when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdoubleword31+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



Format: LWPC rt, offset

microMIPS32 Release 6

Purpose: Load Word PC-relative

To load a word from memory as a signed value, using a PC-relative address.

Description: $GPR[rt] \leftarrow memory[PC \& \sim 0x3 + sign_extend(offset \ll 2)]$

The offset is shifted left by 2 bits, sign-extended, and added to the address of the LWPC instruction.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*.

Restrictions:

LWPC is naturally aligned, by specification.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Operation

```
vAddr ← ( PC & ~0x3 + sign_extend(offset)<<2)
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
memword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
GPR[rt] ← sign_extend(memword)
```


Exceptions:

TLB Refill, TLB Invalid, TLB Read Inhibit, Bus Error, Address Error, Watch

Programming Note

The Release 6 PC-relative loads (LWPC, LWUPC, LDPC) are considered data references.

For the purposes of watchpoints (provided by the CP0 *WatchHi* and *WatchLo* registers) and EJTAG breakpoints, the PC-relative reference is considered to be a data reference rather than an instruction reference. That is, the watchpoint or breakpoint is triggered only if enabled for data references.



Format: LWUPC *rt*, *offset*

microMIPS64 Release 6

Purpose: Load Word Unsigned PC-relative

To load a word from memory as an unsigned value, using a PC-relative address.

Description: $GPR[rt] \leftarrow memory[PC + sign_extend(offset \ll 2)]$

The 19-bit offset is shifted left by 2 bits, sign-extended, and added to the address of the LWUPC instruction.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, zero-extended to the GPR register length if necessary, and placed in GPR *rt*.

Restrictions:

LWUPC is naturally aligned, by specification.

Availability and Compatibility:

This instruction is introduced by and required as of MIPS64 Release 6.

Operation

```
vAddr ← ( PC + sign_extend(offset) << 2 )
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
memword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
GPR[rt] ← zero_extend(memword)
```


Exceptions:

TLB Refill, TLB Invalid, TLB Read Inhibit, Bus Error, Address Error, Watch

Programming Note

The Release 6 PC-relative loads (LWPC, LWUPC, LDPC) are considered data references.

For the purposes of watchpoints (provided by the CP0 *WatchHi* and *WatchLo* registers) and EJTAG breakpoints, the PC-relative reference is considered to be a data reference rather than an instruction reference. That is, the watchpoint or breakpoint is triggered only if enabled for data references.



Format: LWU *rt*, *offset*(*base*)

microMIPS64

Purpose: Load Word Unsigned

To load a word from memory as an unsigned value.

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support because the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdoubleword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← 032 || memdoubleword31+8*byte..8*byte

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Reserved Instruction, Watch

31	26	25	21	20	16	15	11	10	9	8	0
POOL32F 010101	ft		fs		fd		fmt				MADDF 110111000
POOL32F 010101	ft		fs		fd		fmt				MSUBF 111111000
6	5		5		5		2				9

Format: MADDF.fmt MSUBF.fmt
MADDF.S fd, fs, ft
MADDF.D fd, fs, ft
MSUBF.S fd, fs, ft
MSUBF.D fd, fs, ft

microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6

Purpose: Floating Point Fused Multiply Add, Floating Point Fused Multiply Subtract

MADDF.fmt: To perform a fused multiply-add of FP values.

MSUBF.fmt: To perform a fused multiply-subtract of FP values.

Description:

MADDF.fmt: $FPR[fd] \leftarrow FPR[fd] + (FPR[fs] \times FPR[ft])$
MSUBF.fmt: $FPR[fd] \leftarrow FPR[fd] - (FPR[fs] \times FPR[ft])$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The intermediate product is calculated to infinite precision. The product is added to the value in FPR *fd*. The result sum is calculated to infinite precision, rounded according to the current rounding mode in *FCSR*, and placed into FPR *fd*. The operands and result are values in format *fmt*.

(For MSUBF.fmt, the product is subtracted from the value in FPR *fd*.)

Cause bits are ORed into the *Flag* bits if no exception is taken.

Restrictions:

None

Availability and Compatibility:

MADDF.fmt and MSUBF.fmt are required in Release 6.

MADDF.fmt and MSUBF.fmt are not available in architectures pre-Release 6.

The fused multiply add instructions, MADDF.fmt and MSUBF.fmt, replace pre-Release 6 instructions such as MADD.fmt, MSUB.fmt, NMADD.fmt, and NMSUB.fmt. The replaced instructions were unfused multiply-add, with an intermediate rounding.

Release 6 MSUBF.fmt, $fd \leftarrow fd - fs \times ft$, corresponds more closely to pre-Release 6 NMADD.fmt, $fd \leftarrow fr - fs \times ft$, than to pre-Release 6 MSUB.fmt, $fd \leftarrow fs \times ft - fr$.

FPU scalar MADDF.fmt corresponds to MSA vector MADD.df.

FPU scalar MSUBF.fmt corresponds to MSA vector MSUB.df.

Operation:

```
if not IsCoprorocessorEnabled(1)
    then SignalException(CoprorocessorUnusable, 1) endif
if not IsFloatingPointImplemented(fmt)
    then SignalException(ReservedInstruction) endif
```

```

vfr ← ValueFPR(fr, fmt)
vfs ← ValueFPR(fs, fmt)
vfd ← ValueFPR(fd, fmt)
MADDF.fmt: vinf ← vfd +∞ (vfs *∞ vft)
MADDF.fmt: vinf ← vfd -∞ (vfs *∞ vft)
StoreFPR(fd, fmt, vinf)

```

Special Considerations:

The fused multiply-add computation is performed in infinite precision, and signals Inexact, Overflow, or Underflow if and only if the final result differs from the infinite precision result in the appropriate manner.

Like most FPU computational instructions, if the flush-subnormals-to-zero mode, FCSR.FS=1, then subnormals are flushed before beginning the fused-multiply-add computation, and Inexact may be signaled.


I.e. Inexact may be signaled both by input flushing and/or by the fused-multiply-add: the conditions or ORed.

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow, Underflow



Format: MAX.fmt MIN.fmt MAXA.fmt MINA.fmt

MAX.S fd, fs, ft

MAX.D fd, fs, ft

MAXA.S fd, fs, ft

MAXA.D fd, fs, ft

MIN.S fd, fs, ft

MIN.D fd, fs, ft

MINA.S fd, fs, ft

MINA.D fd, fs, ft

microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6
microMIPS32 Release 6

Purpose: Scalar Floating-Point Max/Min/maxNumMag/minNumMag

Scalar Floating-Point Maximum

Scalar Floating-Point Minimum

Scalar Floating-Point argument with Maximum Absolute Value

Scalar Floating-Point argument with Minimum Absolute Value

Description:

MAX.fmt: $FPR[fd] \leftarrow \maxNum(FPR[fs], FPR[ft])$

MIN.fmt: $FPR[fd] \leftarrow \minNum(FPR[fs], FPR[ft])$

MAXA.fmt: $FPR[fd] \leftarrow \maxNumMag(FPR[fs], FPR[ft])$

MINA.fmt: $FPR[fd] \leftarrow \minNumMag(FPR[fs], FPR[ft])$

MAX.fmt writes the maximum value of the inputs *fs* and *ft* to the destination *fd*.

MIN.fmt writes the minimum value of the inputs *fs* and *ft* to the destination *fd*.

MAXA.fmt takes input arguments *fs* and *ft* and writes the argument with the maximum absolute value to the destination *fd*.

MINA.fmt takes input arguments *fs* and *ft* and writes the argument with the minimum absolute value to the destination *fd*.

The instructions MAX.fmt/MIN.fmt/MAXA.fmt/MINA.fmt correspond to the IEEE 754-2008 operations *maxNum*/

minNum/maxNumMag/minNumMag.

- MAX.fmt corresponds to the IEEE 754-2008 operation maxNum.
- MIN.fmt corresponds to the IEEE 754-2008 operation minNum.
- MAXA.fmt corresponds to the IEEE 754-2008 operation maxNumMag.
- MINA.fmt corresponds to the IEEE 754-2008 operation minNumMag.

Numbers are preferred to NaNs: if one input is a NaN, but not both, the value of the numeric input is returned. If both are NaNs, the NaN in `fs` is returned.¹

The scalar FPU instructions MAX.fmt/MIN.fmt/MAXA.fmt/MINA.fmt correspond to the MSA instructions FMAX.df/FMIN.df/FMAXA.df/FMINA.df.

- Scalar FPU instruction MAX.fmt corresponds to the MSA vector instruction FMAX.df.
- Scalar FPU instruction MIN.fmt corresponds to the MSA vector instruction FMIN.df.
- Scalar FPU instruction MAXA.fmt corresponds to the MSA vector instruction FMAX_A.df.
- Scalar FPU instruction MINA.fmt corresponds to the MSA vector instruction FMIN_A.df.

Restrictions:

Data-dependent exceptions are possible as specified by the IEEE Standard for Floating-Point Arithmetic 754TM-2008. See also the section “Special Cases”, below.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Operation:

```

if not IsCoprocesorEnabled(1)
    then SignalException(CoprocesorUnusable, 1) endif
if not IsFloatingPointImplemented(fmt)
    then SignalException(ReservedInstruction) endif

v1 ← ValueFPR(fs, fmt)
v2 ← ValueFPR(ft, fmt)

if SNaN(v1) or SNaN(v2) then
    then SignalException(InvalidOperand) endif

if NaN(v1) and NaN(v2) then
    ftmp ← v1
elseif NaN(v1) then
    ftmp ← v2
elseif NaN(v2) then
    ftmp ← v1
else
    case instruction of

```

1. IEEE standard 754-2008 allows either input to be chosen if both inputs are NaNs. Release 6 specifies that the first input must be propagated.


```

    FMAX.fmt:  ftmp ← MaxFP.fmt (ValueFPR(fs,fmt),ValueFPR(ft,fmt))
    FMIN.fmt:  ftmp ← MinFP.fmt (ValueFPR(fs,fmt),ValueFPR(ft,fmt))
    FMAXA.fmt: ftmp ← MaxAbsoluteFP.fmt (ValueFPR(fs,fmt),ValueFPR(ft,fmt))
    FMINA.fmt: ftmp ← MinAbsoluteFP.fmt (ValueFPR(fs,fmt),ValueFPR(ft,fmt))
  end case
endif

StoreFPR (fd, fmt, ftmp)
/* end of instruction */

function MaxFP(tt, ts, n)
  /* Returns the largest argument. */
endfunction MaxFP

function MinFP(tt, ts, n)
  /* Returns the smallest argument. */
endfunction MinFP

function MaxAbsoluteFP(tt, ts, n)
  /* Returns the argument with largest absolute value.
     For equal absolute values, returns the largest argument.*/
endfunction MaxAbsoluteFP

function MinAbsoluteFP(tt, ts, n)
  /* Returns the argument with smallest absolute value.
     For equal absolute values, returns the smallest argument.*/
endfunction MinAbsoluteFP

function NaN(tt, ts, n)
  /* Returns true if the value is a NaN */
  return SNaN(value) or QNaN(value)
endfunction NaN

```

Table 7.24 Special Cases for FP MAX, MIN, MAXA, MINA

Operand		Other	Release 6 Instructions			
fs	ft		MAX	MIN	MAXA	MINA
-0.0	0.0		0.0	-0.0	0.0	-0.0
0.0	-0.0					
QNaN	#		#	#	#	#
#	QNaN					
QNaN1	QNaN2	Release 6	QNaN1	QNaN1	QNaN1	QNaN1
		IEEE 754 2008	Arbitrary choice. Not allowed to clear sign bit.			

Table 7.24 Special Cases for FP MAX, MIN, MAXA, MINA


Operand		Other	Release 6 Instructions			
fs	ft		MAX	MIN	MAXA	MINA
Either or both operands SNaN		Invalid Operation exception enabled	Signal Invalid Operation Exception. Destination not written.			
		... disabled	Treat as if the SNaN were a QNaN (do not quieten the result).			

Exceptions:

Coproprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation



Format: MFC0 rt, rs
MFC0 rt, rs, sel

microMIPS
microMIPS

Purpose: Move from Coprocessor 0

To move the contents of a coprocessor 0 register to a general register.

Description: $GPR[rt] \leftarrow CPR[0,rs,sel]$

The contents of the coprocessor 0 register specified by the combination of *rs* and *sel* are sign-extended and loaded into general register *rt*. Not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be zero.

When the coprocessor 0 register specified is the *EntryLo0* or the *EntryLo1* register, the RI/XI fields are moved to bits 31:30 of the destination register. This feature supports MIPS32 backward compatibility on a MIPS64 system.

Restrictions:

Pre-Release 6: The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rs* and *sel*.

Release 6: Reading a reserved register or a register that is not implemented for the current core configuration returns 0.

Operation:


```

reg = rs
if IsCoprocessorRegisterImplemented(0, reg, sel) then
    data ← CPR[0, reg, sel]
    if (reg,sel = EntryLo1 or reg,sel = EntryLo0) then
        GPR[rt]29..0 ← data29..0
        GPR[rt]31 ← data63
        GPR[rt]30 ← data62
        GPR[rt]63..32 ← sign_extend(data63)
    else
        GPR[rt] ← sign_extend(data)
    endif
else
    if ArchitectureRevision() ≥ 6 then
        GPR[rt] ← 0
    else
        UNDEFINED
    endif
endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MFC1 rt, fs

microMIPS

Purpose: Move Word From Floating Point

To copy a word from an FPU (CP1) general register to a GPR.

Description: $GPR[rt] \leftarrow FPR[fs]$

The contents of FPR *fs* are sign-extended and loaded into general register *rt*.

Restrictions:

Operation:


```
data ← ValueFPR(fs, UNINTERPRETED_WORD) 31..0
GPR[rt] ← sign_extend(data)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Historical Information:

For MIPS I, MIPS II, and MIPS III the contents of GPR *rt* are **UNPREDICTABLE** for the instruction immediately following MFC1.



Format: MFC2 rt, Impl

microMIPS

The syntax shown above is an example using MFC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word From Coprocessor 2

To copy a word from a COP2 general register to a GPR.

Description: $\text{GPR}[\text{rt}] \leftarrow \text{CP2CPR}[\text{Impl}]$

The contents of the coprocessor 2 register denoted by the *Impl* field are sign-extended and placed into general register *rt*. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if the *Impl* field specifies a coprocessor 2 register that does not exist.

Operation:

```
data ← CP2CPR[Impl]31..0
GPR[rt] ← sign_extend(data)
```

Exceptions:

Coprocessor Unusable

31	26	25	21	20	16	15	14	13	11	10	6	5	0
POOL32A 000000	rt		rs		00		sel		MFHC0 00011		POOL32P 110100		
6	5		5		3		5		6				

Format: MFHC0 rt, rs
MFHC0 rt, rs, sel

microMIPS Release 5
microMIPS Release 5

Purpose: Move from High Coprocessor 0

To move the contents of the upper 32 bits of a Coprocessor 0 register, extended by 32-bits, to a general register.

Description: $GPR[rt] \leftarrow CPR[0,rs,sel][63:32]$

The contents of the Coprocessor 0 register specified by the combination of *rs* and *sel* are sign-extended and loaded into general register *rt*. Not all Coprocessor 0 registers support the *sel* field, and in those instances, the *sel* field must be zero.

When the Coprocessor 0 register specified is the *EntryLo0* or the *EntryLo1* register, MFHC0 must undo the effects of MTHC0, that is, bits 31:30 of the register must be returned as bits 1:0 of the GPR, and bits 32 and those of greater significance must be left-shifted by two and written to bits 31:2 of the GPR. This is because the RI and XI bits are repositioned on the write from GPR to *EntryLo0* or the *EntryLo1*.

Restrictions:

Pre-Release 6: The results are **UNDEFINED** if Coprocessor 0 does not contain a register as specified by *rs* and *sel*, or the register exists but is not extended by 32-bits, or the register is extended for XPA, but XPA is not supported or enabled.

Release 6: Reading the high part of a register that is reserved, not implemented for the current core configuration, or that is not extended beyond 32 bits returns 0.

Availability and Compatibility:

This feature supports MIPS32 backward-compatibility on MIPS64 implementations.

Operation:

PABITS is the total number of physical address bits implemented. PABITS is defined in the descriptions of *EntryLo0* and *EntryLo1*.

```


if Config5MVH = 0 then SignalException(ReservedInstruction) endif
reg ← rs
if IsCoprocessorRegisterImplemented(0, reg, sel) and
   IsCoprocessorRegisterExtended(0, reg, sel) then
    data ← CPR[0, reg, sel]
    if (reg,sel = EntryLo1 or reg,sel = EntryLo0) then
        if (Config3LPA = 1 and PageGrainELPA = 1) then // PABITS > 36
            GPR[rt]31..0 ← data61..30
            GPR[rt]63..32 ← (data61)32 // sign-extend
        else
            GPR[rt] ← 0
        endif
    else
        GPR[rt] ← sign_extend(data63..32)
    endif
else
    if ArchitectureRevision() ≥ 6 then

```

```
        GPR[rt] ← 0
    else
        UNDEFINED
    endif
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MFHC1 rt, fs

microMIPS

Purpose: Move Word From High Half of Floating Point Register

To copy a word from the high half of an FPU (CP1) general register to a GPR.

Description: $GPR[rt] \leftarrow \text{sign_extend}(FPR[fs]_{63..32})$

The contents of the high word of FPR *fs* are sign-extended and loaded into general register *rt*. This instruction is primarily intended to support 64-bit floating point units on a 32-bit CPU, but the semantics of the instruction are defined for all cases.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.


The results are **UNPREDICTABLE** if $Status_{FR} = 0$ and *fs* is odd.

Operation:

```
data ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)_{63..32}
GPR[rt] ← sign_extend(data)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MFHC2 *rt*, *Impl*

microMIPS

The syntax shown above is an example using MFHC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word From High Half of Coprocessor 2 Register

To copy a word from the high half of a COP2 general register to a GPR.

Description: $GPR[rt] \leftarrow \text{sign_extend}(CP2CPR[Impl]_{63..32})$

The contents of the high word of the coprocessor 2 register denoted by the *Impl* field are sign-extended and placed into GPR *rt*. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if the *Impl* field specifies a coprocessor 2 register that does not exist, or if that register is not 64 bits wide.


In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

```
data ← CP2CPR[Impl]63..32
GPR[rt] ← sign_extend(data)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MOV.fmt
 MOV.S ft, fs
 MOV.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Move

To move an FP value between FPRs.

Description: $FPR[ft] \leftarrow FPR[fs]$

The value in FPR *fs* is placed into FPR *ft*. The source and destination are values in format *fmt*. In paired-single format, both the halves of the pair are copied to *ft*.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Availability and Compatibility:

MOV.PS has been removed in Release 6.

Operation:

StoreFPR(ft, fmt, ValueFPR(fs, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation

31	26	25	21	20	16	15	14	13	11	10	6	5	0
POOL32A 000000	rt					rs		00	sel		MTC0 01011		POOL32AXf 111100
6	5					5		2	3		5		6

Format: MTC0 rt, rs
MTC0 rt, rs, sel

microMIPS
microMIPS

Purpose: Move to Coprocessor 0

To move the contents of a general register to a coprocessor 0 register.

Description: $CPR[0, rs, sel] \leftarrow GPR[rt]$

The contents of general register *rt* are loaded into the coprocessor 0 register specified by the combination of *rs* and *sel*. Not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be set to zero.

When the CP0 destination register specified is the *EntryLo0* or the *EntryLo1* register, bits 31:30 appear in the RI/XI fields of the destination register. This feature supports MIPS32 backward compatibility on a MIPS64 implementation.

Restrictions:

Pre-Release 6: The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rs* and *sel*.

Release 6: Writes to a register that is reserved or not defined for the current core configuration are ignored.

Operation:


```

data ← GPR[rt]
reg ← rs
if IsCoproprocessorRegisterImplemented (0, reg, sel) then
    if (reg, sel = EntryLo1 or EntryLo0) then
        CPR[0, reg, sel]29..0 ← data29..0
        CPR[0, reg, sel]63 ← data31
        CPR[0, reg, sel]62 ← data30
        CPR[0, reg, sel]61:30 ← 032
    elseif (Width(CPR[0, reg, sel]) = 64) then
        CPR[0, reg, sel] ← data
    else
        CPR[0, reg, sel] ← data31..0
    endif
else
    if ArchitectureRevision() ≥ 6 then
        // nop (no exceptions, coprocessor state not modified)
    else
        UNDEFINED
    endif
endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MTC1 rt, fs

microMIPS

Purpose: Move Word to Floating Point

To copy a word from a GPR *rt* to an FPU (CP1) general register.

Description: $FPR[fs] \leftarrow GPR[rt]$

The low word in GPR *rt* is placed into the low word of FPR *fs*. If FPRs are 64 bits wide, bits 63..32 of FPR *fs* become **UNPREDICTABLE**.

Restrictions:

Operation:


```
data ← GPR[rt]31..0
StoreFPR(fs, UNINTERPRETED_WORD, data)
```

Exceptions:

Coprocessor Unusable

Historical Information:

For MIPS I, MIPS II, and MIPS III the value of FPR *fs* is **UNPREDICTABLE** for the instruction immediately following MTC1.



Format: MTC2 *rt*, *Impl*

microMIPS

The syntax shown above is an example using MTC1 as a model. The specific syntax is implementation-dependent.

Purpose: Move Word to Coprocessor 2

To copy a word from a GPR to a COP2 general register.

Description: $CP2CPR[Impl] \leftarrow GPR[rt]$

The low word in GPR *rt* is placed into the low word of a Coprocessor 2 general register denoted by the *Impl* field. If Coprocessor 2 general registers are 64 bits wide; bits 63..32 of the register denoted by the *Impl* field become **UNPREDICTABLE**. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if the *Impl* field specifies a Coprocessor 2 register that does not exist.

Operation:

```
data ← GPR[rt]31..0
CP2CPR[Impl] ← data
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

31	26	25	21	20	16	15	14	13	11	10	6	5	0
POOL32A 000000	rt	rs	00	sel	MTHC0 01011	POOL32P 110100							
6	5	5	2	3	5	6							

Format: MTHC0 rt, rs
MTHC0 rt, rs, sel

microMIPS Release 5
microMIPS Release 5

Purpose: Move to High Coprocessor 0

To copy a word from a GPR to the upper 32 bits of a CP0 general register that has been extended by 32 bits.

Description: $\text{CPR}[0, \text{rs}, \text{sel}][63:32] \leftarrow \text{GPR}[\text{rt}]$

The contents of general register *rt* are loaded into the Coprocessor 0 register specified by the combination of *rs* and *sel*. Not all Coprocessor 0 registers support the *sel* field; the *sel* field must be set to zero.

When the Coprocessor 0 destination register specified is the *EntryLo0* or *EntryLo1* register, bits 1:0 of the GPR appear at bits 31:30 of *EntryLo0* or *EntryLo1*. This is to compensate for *Rl* and *Xl*, which were shifted to bits 63:62 by MTC0 to *EntryLo0* or *EntryLo1*. If *Rl/Xl* are not supported, the shift must still occur, but an MFHC0 instruction returns 0s for these two fields. The GPR is right-shifted by two to vacate the lower two bits, and two 0s are shifted in from the left. The result is written to the upper 32 bits of MIPS64 *EntryLo0* or *EntryLo1*, excluding *Rl/Xl*, which were placed in bits 63:62, that is, the write must appear atomic, as if both MTC0 and MTHC0 occurred together.

This feature supports MIPS32 backward compatibility of MIPS64 systems.

Restrictions:

Pre-Release 6: The results are **UNDEFINED** if Coprocessor 0 does not contain a register as specified by *rs* and *sel*, or if the register exists but is not extended by 32 bits, or the register is extended for XPA, but XPA is not supported or enabled.

Release 6: A write to the high part of a register that is reserved, not implemented for the current core, or that is not extended beyond 32 bits is ignored.

In a 64-bit processor, the MTHC0 instruction writes only the lower 32 bits of register *rt* into the upper 32 bits of the Coprocessor register specified by *rd* and *sel* if the register is extended as defined by `IsCoprocessorRegisterExtended()`. The registers extended by Release 5 are those required for the XPA feature. Release 6 extends *WatchHi* to support *MemoryMapID*. These registers are identical to the same registers in the MIPS64 Architecture, other than *EntryLo0* and *EntryLo1*.

Operation:

```


if Config5MVH = 0 then SignalException(ReservedInstruction) endif
data ← GPR[rt]
reg ← rs
if IsCoprocessorRegisterImplemented (0, reg, sel) and
   IsCoprocessorRegisterExtended (0, reg, sel) then
    if (reg,sel = EntryLo1 or reg,sel = EntryLo0) then
        if (Config3LPA = 1 and PageGrainELPA = 1) then // PABITS > 36
            CPR[0,reg,sel]31..30 ← data1..0
            CPR[0,reg,sel]61:32 ← data31..2 and ((1<<(PABITS-36))-1)
            CPR[0,reg,sel]63:62 ← 02
        else
            CPR[0, reg, sel][63:32] ← data31..0
        endif
    else
        CPR[0, reg, sel][63:32] ← data31..0
    endif
else
    if ArchitectureRevision() ≥ 6 then

```

```
        // nop (no exceptions, coprocessor state not modified)
    else
        UNDEFINED
    endif
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: MTHC1 rt, fs

microMIPS

Purpose: Move Word to High Half of Floating Point Register

To copy a word from a GPR to the high half of an FPU (CP1) general register.

Description: $\text{FPR}[\text{fs}]_{63..32} \leftarrow \text{GPR}[\text{rt}]_{31..0}$

The low word in GPR *rt* is placed into the high word of FPR *fs*. This instruction is primarily intended to support 64-bit floating point units on a 32-bit CPU, but the semantics of the instruction are defined for all cases.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

The results are **UNPREDICTABLE** if $\text{Status}_{FR} = 0$ and *fs* is odd.

Operation:

```
newdata ← GPR[rt]31..0
olddata ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)31..0
StoreFPR(fs, UNINTERPRETED_DOUBLEWORD, newdata || olddata)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Programming Notes

When paired with MTC1 to write a value to a 64-bit FPR, the MTC1 must be executed first, followed by the MTHC1. This is because of the semantic definition of MTC1, which is not aware that software is using an MTHC1 instruction to complete the operation, and sets the upper half of the 64-bit FPR to an **UNPREDICTABLE** value.

31	26	25	21	20	16	15	6	5	0
POOL32A 000000						rt		Impl	
						MTHC2 1001110100		POOL32AXf 111100	
6						5		5	
						10		6	

Format: MTHC2 rt, Impl

microMIPS

The syntax shown above is an example using MTHC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word to High Half of Coprocessor 2 Register

To copy a word from a GPR to the high half of a COP2 general register.

Description: $CP2CPR[Impl]_{63..32} \leftarrow GPR[rt]_{31..0}$

The low word in GPR *rt* is placed into the high word of coprocessor 2 general register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if the *Impl* field specifies a coprocessor 2 register that does not exist, or if that register is not 64 bits wide.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

```
data ← GPR[rt]31..0
CP2CPR[Impl] ← data || CPR[2,rd,sel]31..0
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Programming Notes

When paired with MTC2 to write a value to a 64-bit CPR, the MTC2 must be executed first, followed by the MTHC2. This is because of the semantic definition of MTC2, which is not aware that software is using an MTHC2 instruction to complete the operation, and sets the upper half of the 64-bit CPR to an **UNPREDICTABLE** value.

31	26	25	21	20	16	15	11	10	9	0	
POOL32A 000000	rt		rs		rd		0		MUL 0000011000		
POOL32A 000000	rt		rs		rd		0		MUH 0001011000		
POOL32A 000000	rt		rs		rd		0		MULU 0010011000		
POOL32A 000000	rt		rs		rd		0		MUHU 0011011000		
6	5		5		5		1		10		
31	26	25	21	20	16	15	11	10	9	8	0
POOL32S 010110	rt		rs		rd		00		DMUL 000011000		
POOL32S 010110	rt		rs		rd		00		DMUH 001011000		
POOL32S 010110	rt		rs		rd		00		DMULU 010011000		
POOL32S 010110	rt		rs		rd		00		DMUHU 011011000		
6	5		5		5		2		9		

Format: MUL MUH MULU MUHU DMUL DMUH DMULU DMUHU

MUL rd,rs,rt
 MUH rd,rs,rt
 MULU rd,rs,rt
 MUHU rd,rs,rt
 DMUL rd,rs,rt
 DMUH rd,rs,rt
 DMULU rd,rs,rt
 DMUHU rd,rs,rt

microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6
 microMIPS64 Release 6

Purpose: Multiply Integers (with result to GPR)

MUL: Multiply Words Signed, Low Word
 MUH: Multiply Words Signed, High Word
 MULU: Multiply Words Unsigned, Low Word
 MUHU: Multiply Words Unsigned, High Word
 DMUL: Multiply Doublewords Signed, Low Doubleword
 DMUH: Multiply Doublewords Signed, High Doubleword
 DMULU: Multiply Doublewords Unsigned, Low Doubleword
 DMUHU: Multiply Doublewords Unsigned, High Doubleword

Description:

MUL: GPR[rd] ← sign_extend.32(lo_word(multiply.signed(GPR[rs] × GPR[rt])))
 MUH: GPR[rd] ← sign_extend.32(hi_word(multiply.signed(GPR[rs] × GPR[rt])))
 MULU: GPR[rd] ← sign_extend.32(lo_word(multiply.unsigned(GPR[rs] × GPR[rt])))
 MUHU: GPR[rd] ← sign_extend.32(hi_word(multiply.unsigned(GPR[rs] × GPR[rt])))
 DMUL: GPR[rd] ← lo_doubleword(multiply.signed(GPR[rs] × GPR[rt]))
 DMUH: GPR[rd] ← hi_doubleword(multiply.signed(GPR[rs] × GPR[rt]))

```
DMULU: GPR[rd] ← lo_doubleword( multiply.unsigned( GPR[rs] × GPR[rt] ) )
DMUHU: GPR[rd] ← hi_doubleword( multiply.unsigned( GPR[rs] × GPR[rt] ) )
```

The Release 6 multiply instructions multiply the operands in GPR[rs] and GPR[rt], and place the specified high or low part of the result, of the same width, in GPR[rd].

MUL performs a signed 32-bit integer multiplication, and places the low 32 bits of the result in the destination register.

MUH performs a signed 32-bit integer multiplication, and places the high 32 bits of the result in the destination register.

MULU performs an unsigned 32-bit integer multiplication, and places the low 32 bits of the result in the destination register.

MUHU performs an unsigned 32-bit integer multiplication, and places the high 32 bits of the result in the destination register.

DMUL performs a signed 64-bit integer multiplication, and places the low 64 bits of the result in the destination register.

DMUH performs a signed 64-bit integer multiplication, and places the high 64 bits of the result in the destination register.

DMULU performs an unsigned 64-bit integer multiplication, and places the low 64 bits of the result in the destination register.

DMUHU performs an unsigned 64-bit integer multiplication, and places the high 64 bits of the result in the destination register.

Restrictions:

On a 64-bit CPU, MUH is UNPREDICTABLE if its inputs are not signed extended 32-bit integers.

MUL behaves correctly even if its inputs are not sign extended 32-bit integers. Bits 32-63 of its inputs do not affect the result.

On a 64-bit CPU, MUHU is UNPREDICTABLE if its inputs are not zero or sign extended 32-bit integers.

MULU behaves correctly even if its inputs are not zero or sign extended 32-bit integers. Bits 32-63 of its inputs do not affect the result.

On a 64-bit CPU, the 32-bit multiplications, both signed and unsigned, sign extend the result as if it is a 32-bit signed integer.

DMUL DMUH DMULU DMUHU: Reserved Instruction exception if 64-bit instructions are not enabled.

Availability and Compatibility:

These instructions are introduced by and required as of Release 6.

Programming Notes:

The low half of the integer multiplication result is identical for signed and unsigned. Nevertheless, there are distinct instructions MUL MULU DMUL DMULU. Implementations may choose to optimize a multiply that produces the low half followed by a multiply that produces the upper half. Programmers are recommended to use matching lower and upper half multiplications.

The Release 6 MUL instruction has the same opcode mnemonic as the pre-Release 6 MUL instruction. The semantics of these instructions are almost identical: both produce the low 32-bits of the 32×32=64 product; but the pre-Release 6 MUL is unpredictable if its inputs are not properly sign extended 32-bit values on a 64 bit machine, and is defined to render the HI and LO registers unpredictable, whereas the Release 6 version ignores bits 32-63 of the input, and there are no HI/LO registers in Release 6 to be affected.

Operation:

```

MUH: if NotWordValue(GPR[rs]) then UNPREDICTABLE endif
MUH: if NotWordValue(GPR[rt]) then UNPREDICTABLE endif
MUHU: if NotWordValue(GPR[rs]) then UNPREDICTABLE endif
MUHU: if NotWordValue(GPR[rt]) then UNPREDICTABLE endif

/* recommended implementation: ignore bits 32-63 for MUL, MUH, MULU, MUHU */

MUL, MUH:
    s1 ← signed_word(GPR[rs])
    s2 ← signed_word(GPR[rt])
MULU, MUHU:
    s1 ← unsigned_word(GPR[rs])
    s2 ← unsigned_word(GPR[rt])
DMUL, DMUH:
    s1 ← signed_doubleword(GPR[rs])
    s2 ← signed_doubleword(GPR[rt])
DMULU, DMUHU:
    s1 ← unsigned_doubleword(GPR[rs])
    s2 ← unsigned_doubleword(GPR[rt])

product ← s1 × s2      /* product is twice the width of sources */


MUL:  GPR[rd] ← sign_extend.32( lo_word( product ) )
MUH:  GPR[rd] ← sign_extend.32( hi_word( product ) )
MULU: GPR[rd] ← sign_extend.32( lo_word( product ) )
MUHU: GPR[rd] ← sign_extend.32( hi_word( product ) )
DMUL: GPR[rd] ← lo_doubleword( product )
DMUH: GPR[rd] ← hi_doubleword( product )
DMULU: GPR[rd] ← lo_doubleword( product )
DMUHU: GPR[rd] ← hi_doubleword( product )

```

Exceptions:

MUL MUH MULU MUHU: None

DMUL DMUH DMULU DMUHU: Reserved Instruction



Format: MUL.fmt
MUL.S fd, fs, ft
MUL.D fd, fs, ft

microMIPS
microMIPS

Purpose: Floating Point Multiply

To multiply FP values.

Description: $FPR[fd] \leftarrow FPR[fs] \times FPR[ft]$

The value in FPR *fs* is multiplied by the value in FPR *ft*. The result is calculated to infinite precision, rounded according to the current rounding mode in *FCSR*, and placed into FPR *fd*. The operands and result are values in format *fmt*.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

Availability and Compatibility:

MUL.PS has been removed in Release 6.

Operation:


StoreFPR (fd, fmt, ValueFPR(fs, fmt) ∞_{fmt} ValueFPR(ft, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow, Underflow



Format: NEG.fmt
 NEG.S ft, fs
 NEG.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Negate

To negate an FP value.

Description: $FPR[ft] \leftarrow -FPR[fs]$

The value in FPR *fs* is negated and placed into FPR *ft*. The value is negated by changing the sign bit value. The operand and result are values in format *fmt*.

If $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$ then this operation is arithmetic. For this case, any NaN operand signals invalid operation.

If $FCSR_{ABS2008}=1$ then this operation is non-arithmetic. For this case, both regular floating point numbers and NaN values are treated alike, only the sign bit is affected by this instruction. No IEEE 754 exception can be generated for this case.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**. The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Availability and Compatibility:

NEG.PS has been removed in Release 6.

Operation:

`StoreFPR(ft, fmt, Negate(ValueFPR(fs, fmt)))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation

31	26	25	21	20	16	15	11	10	5	0
POOL32A 000000	0 00000	0 00000	0 00000	0 00000	0 00000	0 00000	0 00000	0 00000	SLL 000000	
6	5	5	5	5	5	5	5	5	6	

Format: NOP

Assembly Idiom microMIPS

Purpose: No Operation

To perform no operation.

Description:

NOP is the assembly idiom used to denote no operation. The actual instruction is interpreted by the hardware as SLL r0, r0, 0.

Restrictions:

None

Operations:


None

Exceptions:

None

Programming Notes:

The zero instruction word, which represents SLL, r0, r0, 0, is the preferred NOP for software to use and to pad out alignment sequences.



Format: NOR *rd*, *rs*, *rt*

microMIPS

Purpose: Not Or

To do a bitwise logical NOT OR.

Description: $GPR[rd] \leftarrow GPR[rs] \text{ nor } GPR[rt]$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical NOR operation. The result is placed into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ nor } GPR[rt]$

Exceptions:

None



Format: OR *rd*, *rs*, *rt*

microMIPS

Purpose: Or

To do a bitwise logical OR.

Description: $GPR[rd] \leftarrow GPR[rs] \text{ or } GPR[rt]$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical OR operation. The result is placed into GPR *rd*.

Restrictions:


None

Operations:

$GPR[rd] \leftarrow GPR[rs] \text{ or } GPR[rt]$

Exceptions:

None



Format: ORI *rt*, *rs*, *immediate*

microMIPS

Purpose: Or Immediate

To do a bitwise logical OR with a constant.

Description: $GPR[rt] \leftarrow GPR[rs] \text{ or } immediate$

The 16-bit *immediate* is zero-extended to the left and combined with the contents of GPR *rs* in a bitwise logical OR operation. The result is placed into GPR *rt*.

Restrictions:

None

Operations:

$GPR[rt] \leftarrow GPR[rs] \text{ or } zero_extend(immediate)$

Exceptions:

None

31	26	25			6	5	0
POOL32A 000000	0 00000	0 00000	5 00101	0 00000	SLL 000000		
6	5	5	5	5	6		

Format: PAUSE

microMIPS

Purpose: Wait for the LLBit to clear.**Description:**

Locks implemented using the LL/SC (or LLD/SCD) instructions are a common method of synchronization between threads of control. A lock implementation does a load-linked instruction and checks the value returned to determine whether the software lock is set. If it is, the code branches back to retry the load-linked instruction, implementing an active busy-wait sequence. The PAUSE instruction is intended to be placed into the busy-wait sequence to block the instruction stream until such time as the load-linked instruction has a chance to succeed in obtaining the software lock.

The PAUSE instruction is implementation-dependent, but it usually involves descheduling the instruction stream until the LLBit is zero.

- In a single-threaded processor, this may be implemented as a short-term WAIT operation which resumes at the next instruction when the LLBit is zero or on some other external event such as an interrupt.
- On a multi-threaded processor, this may be implemented as a short term YIELD operation which resumes at the next instruction when the LLBit is zero.

In either case, it is assumed that the instruction stream which gives up the software lock does so via a write to the lock variable, which causes the processor to clear the LLBit as seen by this thread of execution.

The encoding of the instruction is such that it is backward compatible with all previous implementations of the architecture. The PAUSE instruction can therefore be placed into existing lock sequences and treated as a NOP by the processor, even if the processor does not implement the PAUSE instruction.

Restrictions:

Pre-Release 6: The operation of the processor is **UNPREDICTABLE** if a PAUSE instruction is executed placed in the delay slot of a branch or jump instruction. This restriction does not apply in Release 6.

Operations:

```

if LLBit ≠ 0 then
    EPC ← PC + 4                /* Resume at the following instruction */
    DescheduleInstructionStream()
endif

```

Exceptions:

None

Programming Notes:

The PAUSE instruction is intended to be inserted into the instruction stream after an LL instruction has set the LLBit and found the software lock set. The program may wait forever if a PAUSE instruction is executed and there is no possibility that the LLBit will ever be cleared.

An example use of the PAUSE instruction is shown below:


```

acquire_lock:
    ll    t0, 0(a0)           /* Read software lock, set hardware lock */
    bnezc t0, acquire_lock_retry: /* Branch if software lock is taken; */
                                /* Release 6 branch */
    addiu t0, t0, 1           /* Set the software lock */
    sc    t0, 0(a0)           /* Try to store the software lock */
    bnezc t0, 10f             /* Branch if lock acquired successfully */
    sync
acquire_lock_retry:
    pause                               /* Wait for LLBIT to clear before retry */
    bc    acquire_lock           /* and retry the operation; Release 6 branch */
10:

    Critical region code

release_lock:
    sync
    sw    zero, 0(a0)          /* Release software lock, clearing LLBIT */
                                /* for any PAUSED waiters */

```



Format: PREF hint,offset(base)

microMIPS

Purpose: Prefetch

To move data between memory and cache.

Description: prefetch_memory(GPR[base] + offset)

PREF adds the signed *offset* to the contents of GPR *base* to form an effective byte address. The *hint* field supplies information about the way that the data is expected to be used.

PREF enables the processor to take some action, typically causing data to be moved to or from the cache, to improve program performance. The action taken for a specific PREF instruction is both system and context dependent. Any action, including doing nothing, is permitted as long as it does not change architecturally visible state or alter the meaning of a program. Implementations are expected either to do nothing, or to take an action that increases the performance of the program. The PrepareForStore function is unique in that it may modify the architecturally visible state.

PREF does not cause addressing-related exceptions, including TLB exceptions. If the address specified would cause an addressing exception, the exception condition is ignored and no data movement occurs. However even if no data is moved, some action that is not architecturally visible, such as writeback of a dirty cache line, can take place.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREF instruction.

PREF neither generates a memory operation nor modifies the state of a cache line for a location with an *uncached* memory access type, whether this type is specified by the address segment (e.g., kseg1), the programmed cacheability and coherency attribute of a segment (e.g., the use of the *K0*, *KU*, or *K23* fields in the *Config* register), or the per-page cacheability and coherency attribute provided by the TLB.

If PREF results in a memory operation, the memory access type and cacheability&coherency attribute used for the operation are determined by the memory access type and cacheability&coherency attribute of the effective address, just as it would be if the memory operation had been caused by a load or store to the effective address.

For a cached location, the expected and useful action for the processor is to prefetch a block of data that includes the effective address. The size of the block and the level of the memory hierarchy it is fetched into are implementation specific.

In coherent multiprocessor implementations, if the effective address uses a coherent Cacheability and Coherency Attribute (CCA), then the instruction causes a coherent memory transaction to occur. This means a prefetch issued on one processor can cause data to be evicted from the cache in another processor.

The PREF instruction and the memory transactions which are sourced by the PREF instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

Table 8.25 Values of *hint* Field for PREF Instruction

Value	Name	Data Use and Desired Prefetch Action
0	load	Use: Prefetched data is expected to be read (not modified). Action: Fetch data as if for a load.

Table 8.25 Values of *hint* Field for PREF Instruction (Continued)

Value	Name	Data Use and Desired Prefetch Action
1	store	Use: Prefetched data is expected to be stored or modified. Action: Fetch data as if for a store.
2	L1 LRU hint	Pre-Release 6: Reserved for Architecture. Release 6: Implementation dependent. This hint code marks the line as LRU in the L1 cache and thus preferred for next eviction. Implementations can choose to writeback and/or invalidate as long as no architectural state is modified.
3	Reserved for Implementation	Pre-Release 6: Reserved for Architecture. Release 6: Available for implementation-dependent use.
4	load_streamed	Use: Prefetched data is expected to be read (not modified) but not reused extensively; it “streams” through cache. Action: Fetch data as if for a load and place it in the cache so that it does not displace data prefetched as “retained.”
5	store_streamed	Use: Prefetched data is expected to be stored or modified but not reused extensively; it “streams” through cache. Action: Fetch data as if for a store and place it in the cache so that it does not displace data prefetched as “retained.”
6	load_retained	Use: Prefetched data is expected to be read (not modified) and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a load and place it in the cache so that it is not displaced by data prefetched as “streamed.”
7	store_retained	Use: Prefetched data is expected to be stored or modified and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a store and place it in the cache so that it is not displaced by data prefetched as “streamed.”
8-15	L2 operation	Pre-Release 6: Reserved for Architecture. Release 6: In the Release 6 architecture, hint codes 8 - 15 are treated the same as hint codes 0 - 7 respectively, but operate on the L2 cache.
16-23	L3 operation	Pre-Release 6: Reserved for Architecture. Release 6: In the Release 6 architecture, hint codes 16 - 23 are treated the same as hint codes 0 - 7 respectively, but operate on the L3 cache.
24	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture - available for implementation-dependent use. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).
25	writeback_invalidate (also known as “nudge”) Reserved for Architecture in Release 6	Pre-Release 6: Use—Data is no longer expected to be used. Action—For a writeback cache, schedule a writeback of any dirty data. At the completion of the writeback, mark the state of any cache lines written back as invalid. If the cache line is not dirty, it is implementation dependent whether the state of the cache line is marked invalid or left unchanged. If the cache line is locked, no action is taken. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).

Table 8.25 Values of *hint* Field for PREF Instruction (Continued)

Value	Name	Data Use and Desired Prefetch Action
26-29	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture—available for implementation-dependent use. Release 6: These hints are not implemented in the Release 6 architecture and generate a Reserved Instruction exception (RI).
30	PrepareForStore Reserved for Architecture in Release 6	Pre-Release 6: Use—Prepare the cache for writing an entire line, without the overhead involved in filling the line from memory. Action—If the reference hits in the cache, no action is taken. If the reference misses in the cache, a line is selected for replacement, any valid and dirty victim is written back to memory, the entire line is filled with zero data, and the state of the line is marked as valid and dirty. Programming Note: Because the cache line is filled with zero data on a cache miss, software must not assume that this action, in and of itself, can be used as a fast bzero-type function. Release 6: This hint is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).
31	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture—available for implementation-dependent use. Release 6: This hint is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).

Restrictions:

None

This instruction does not produce an exception for a misaligned memory address, since it has no memory access size.

Availability and Compatibility:

This instruction has been recoded for Release 6.

Operation:

```
vAddr ← GPR[base] + sign_extend(offset)
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
Prefetch(CCA, pAddr, vAddr, DATA, hint)
```

Exceptions:

Bus Error, Cache Error

Prefetch does not take any TLB-related or address-related exceptions under any circumstances.

Programming Notes:

In the Release 6 architecture, hint codes 2:3, 10:11, 18:19 behave as a NOP if not implemented. Hint codes 24:31 are not implemented (treated as reserved) and always signal a Reserved Instruction exception (RI).

As shown in the instruction drawing above, Release 6 implements a 9-bit offset, whereas all release levels lower than Release 6 of the MIPS architecture implement a 16-bit offset.

Prefetch cannot move data to or from a mapped location unless the translation for that location is present in the TLB. Locations in memory pages that have not been accessed recently may not have translations in the TLB, so prefetch may not be effective for such locations.


Prefetch does not cause addressing exceptions. A prefetch may be used using an address pointer before the validity of

the pointer is determined without worrying about an addressing exception.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREF instruction. Typically, this only occurs in systems which have high-reliability requirements.

Prefetch operations have no effect on cache lines that were previously locked with the CACHE instruction.

Hint field encodings whose function is described as “streamed” or “retained” convey usage intent from software to hardware. Software should not assume that hardware will always prefetch data in an optimal way. If data is to be truly retained, software should use the Cache instruction to lock data into the cache.



Format: `PREFE hint,offset(base)`

microMIPS

Purpose: Prefetch EVA

To move data between user mode virtual address space memory and cache while operating in kernel mode.

Description: `prefetch_memory(GPR[base] + offset)`

PREFE adds the 9-bit signed *offset* to the contents of GPR *base* to form an effective byte address. The *hint* field supplies information about the way that the data is expected to be used.

PREFE enables the processor to take some action, causing data to be moved to or from the cache, to improve program performance. The action taken for a specific PREFE instruction is both system and context dependent. Any action, including doing nothing, is permitted as long as it does not change architecturally visible state or alter the meaning of a program. Implementations are expected either to do nothing, or to take an action that increases the performance of the program. The PrepareForStore function is unique in that it may modify the architecturally visible state.

PREFE does not cause addressing-related exceptions, including TLB exceptions. If the address specified would cause an addressing exception, the exception condition is ignored and no data movement occurs. However even if no data is moved, some action that is not architecturally visible, such as writeback of a dirty cache line, can take place.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREFE instruction.

PREFE neither generates a memory operation nor modifies the state of a cache line for a location with an *uncached* memory access type, whether this type is specified by the address segment (for example, *kseg1*), the programmed cacheability and coherency attribute of a segment (for example, the use of the *K0*, *KU*, or *K23* fields in the *Config* register), or the per-page cacheability and coherency attribute provided by the TLB.

If PREFE results in a memory operation, the memory access type and cacheability & coherency attribute used for the operation are determined by the memory access type and cacheability & coherency attribute of the effective address, just as it would be if the memory operation had been caused by a load or store to the effective address.

For a cached location, the expected and useful action for the processor is to prefetch a block of data that includes the effective address. The size of the block and the level of the memory hierarchy it is fetched into are implementation specific.

In coherent multiprocessor implementations, if the effective address uses a coherent Cacheability and Coherency Attribute (CCA), then the instruction causes a coherent memory transaction to occur. This means a prefetch issued on one processor can cause data to be evicted from the cache in another processor.

The PREFE instruction and the memory transactions which are sourced by the PREFE instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

The PREFE instruction functions in exactly the same fashion as the PREF instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to one.

Table 8.26 Values of *hint* Field for PREFE Instruction

Value	Name	Data Use and Desired Prefetch Action
0	load	Use: Prefetched data is expected to be read (not modified). Action: Fetch data as if for a load.
1	store	Use: Prefetched data is expected to be stored or modified. Action: Fetch data as if for a store.
2	L1 LRU hint	Pre-Release 6: Reserved for Architecture. Release 6: Implementation dependent. This hint code marks the line as LRU in the L1 cache and thus preferred for next eviction. Implementations can choose to writeback and/or invalidate as long as no architectural state is modified.
3	Reserved for Implementation	Pre-Release 6: Reserved for Architecture. Release 6: Available for implementation-dependent use.
4	load_streamed	Use: Prefetched data is expected to be read (not modified) but not reused extensively; it “streams” through cache. Action: Fetch data as if for a load and place it in the cache so that it does not displace data prefetched as “retained.”
5	store_streamed	Use: Prefetched data is expected to be stored or modified but not reused extensively; it “streams” through cache. Action: Fetch data as if for a store and place it in the cache so that it does not displace data prefetched as “retained.”
6	load_retained	Use: Prefetched data is expected to be read (not modified) and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a load and place it in the cache so that it is not displaced by data prefetched as “streamed.”
7	store_retained	Use: Prefetched data is expected to be stored or modified and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a store and place it in the cache so that it is not displaced by data prefetched as “streamed.”
8-15	L2 operation	Pre-Release 6: Reserved for Architecture. Release 6: Hint codes 8 - 15 are treated the same as hint codes 0 - 7 respectively, but operate on the L2 cache.
16-23	L3 operation	Pre-Release 6: Reserved for Architecture. Release 6: Hint codes 16 - 23 are treated the same as hint codes 0 - 7 respectively, but operate on the L3 cache.
24	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture - available for implementation-dependent use. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).

Table 8.26 Values of *hint* Field for PREFE Instruction (Continued)

Value	Name	Data Use and Desired Prefetch Action
25	writeback_invalidate (also known as “nudge”) Reserved for Architecture in Release 6	Pre-Release 6: Use—Data is no longer expected to be used. Action—For a writeback cache, schedule a writeback of any dirty data. At the completion of the writeback, mark the state of any cache lines written back as invalid. If the cache line is not dirty, it is implementation dependent whether the state of the cache line is marked invalid or left unchanged. If the cache line is locked, no action is taken. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).
26-29	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture - available for implementation-dependent use. Release 6: These hint codes are not implemented in the Release 6 architecture and generate a Reserved Instruction exception (RI).
30	PrepareForStore Reserved for Architecture in Release 6	Pre-Release 6: Use—Prepare the cache for writing an entire line, without the overhead involved in filling the line from memory. Action—If the reference hits in the cache, no action is taken. If the reference misses in the cache, a line is selected for replacement, any valid and dirty victim is written back to memory, the entire line is filled with zero data, and the state of the line is marked as valid and dirty. Programming Note: Because the cache line is filled with zero data on a cache miss, software must not assume that this action, in and of itself, can be used as a fast bzero-type function. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).
31	Reserved for Architecture	Pre-Release 6: Unassigned by the Architecture - available for implementation-dependent use. Release 6: This hint code is not implemented in the Release 6 architecture and generates a Reserved Instruction exception (RI).

Restrictions:

Only usable when access to Coprocessor0 is enabled and when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

This instruction does not produce an exception for a misaligned memory address, since it has no memory access size.

Operation:

```

vAddr ← GGPR[base] + sign_extend(offset)
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
Prefetch(CCA, pAddr, vAddr, DATA, hint)

```

Exceptions:

Bus Error, Cache Error, Address Error, Reserved Instruction, Coprocessor Usable

Prefetch does not take any TLB-related or address-related exceptions under any circumstances.

Programming Notes:

In the Release 6 architecture, hint codes 0:23 behave as a NOP and never signal a Reserved Instruction exception (RI). Hint codes 24:31 are not implemented (treated as reserved) and always signal a Reserved Instruction exception (RI).


Prefetch cannot move data to or from a mapped location unless the translation for that location is present in the TLB. Locations in memory pages that have not been accessed recently may not have translations in the TLB, so prefetch may not be effective for such locations.

Prefetch does not cause addressing exceptions. A prefetch may be used using an address pointer before the validity of the pointer is determined without worrying about an addressing exception.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREFE instruction. Typically, this only occurs in systems which have high-reliability requirements.

Prefetch operations have no effect on cache lines that were previously locked with the CACHE instruction.

Hint field encodings whose function is described as “streamed” or “retained” convey usage intent from software to hardware. Software should not assume that hardware will always prefetch data in an optimal way. If data is to be truly retained, software should use the Cache instruction to lock data into the cache.



Format: RDHWR *rt,rs,sel*

microMIPS

Purpose: Read Hardware Register

To move the contents of a hardware register to a general purpose register (GPR) if that operation is enabled by privileged software.

The purpose of this instruction is to give user mode access to specific information that is otherwise only visible in kernel mode.

In Release 6, a *sel* field has been added to allow a register with multiple instances to be read selectively. Specifically it is used for *PerfCtr*.

Description: $GPR[rt] \leftarrow HWR[rs]; GPR[rt] \leftarrow HWR[rs, sel]$

If access is allowed to the specified hardware register, the contents of the register specified by *rs* (optionally *sel* in Release 6) is sign-extended and loaded into general register *rt*. Access control for each register is selected by the bits in the coprocessor 0 *HWREna* register.

The available hardware registers, and the encoding of the *rs* field for each, are shown in [Table 8.27](#).

Table 8.27 RDHWR Register Numbers

Register Number (<i>rs</i> Value)	Mnemonic	Description										
0	CPUNum	Number of the CPU on which the program is currently running. This register provides read access to the coprocessor 0 <i>EBase</i> _{CPUNum} field.										
1	SYNCI_Step	Address step size to be used with the SYNCI instruction, or zero if no caches need be synchronized. See that instruction's description for the use of this value.										
2	CC	High-resolution cycle counter. This register provides read access to the coprocessor 0 <i>Count</i> Register.										
3	CCRes	Resolution of the CC register. This value denotes the number of cycles between update of the register. For example: <table><tr><th>CCRes Value</th><th>Meaning</th></tr><tr><td>1</td><td>CC register increments every CPU cycle</td></tr><tr><td>2</td><td>CC register increments every second CPU cycle</td></tr><tr><td>3</td><td>CC register increments every third CPU cycle</td></tr><tr><td colspan="2">etc.</td></tr></table>	CCRes Value	Meaning	1	CC register increments every CPU cycle	2	CC register increments every second CPU cycle	3	CC register increments every third CPU cycle	etc.	
CCRes Value	Meaning											
1	CC register increments every CPU cycle											
2	CC register increments every second CPU cycle											
3	CC register increments every third CPU cycle											
etc.												
4	PerfCtr	Performance Counter Pair. Even <i>sel</i> selects the <i>Control</i> register, while odd <i>sel</i> selects the <i>Counter</i> register in the pair. The value of <i>sel</i> corresponds to the value of <i>sel</i> used by MFC0 to read the CP0 register.										

Table 8.27 RDHWR Register Numbers

Register Number (<i>rs</i> Value)	Mnemonic	Description
5	XNP	Indicates support for the Release 6 Paired LL/SC family of instructions. If set to 1, the LL/SC family of instructions is not present, otherwise, it is present in the implementation. In absence of hardware support for double-width or extended atomics, user software may emulate the instruction's behavior through other means. See <i>Config5_{XNP}</i> .
6-28		These registers numbers are reserved for future architecture use. Access results in a Reserved Instruction Exception.
29	ULR	User Local Register. This register provides read access to the coprocessor 0 <i>UserLocal</i> register, if it is implemented. In some operating environments, the <i>UserLocal</i> register is a pointer to a thread-specific storage block.
30-31		These register numbers are reserved for implementation-dependent use. If they are not implemented, access results in a Reserved Instruction Exception.

Restrictions:

In implementations of Release 1 of the Architecture, this instruction resulted in a Reserved Instruction Exception.

Access to the specified hardware register is enabled if Coprocessor 0 is enabled, or if the corresponding bit is set in the *HWREna* register. If access is not allowed or the register is not implemented, a Reserved Instruction Exception is signaled.

In Release 6, when the 3-bit *sel* is undefined for use with a specific register number, then a Reserved Instruction Exception is signaled.

Availability and Compatibility:

This instructions has been recoded for Release 6. The instruction supports a *sel* field in Release 6.

Operation:

```

    if ((rs!=4) and (sel==0))
    case rs
        0: temp ← sign_extend(EBaseCPUNum)
        1: temp ← sign_extend(SYNCI_StepSize())
        2: temp ← sign_extend(Count)
        3: temp ← sign_extend(CountResolution())
            if (>=2) // #5 - Release 6
                5: temp ← sign_extend(0x00000001 && Config5XNP) //zero-extend really
            endif
        29: temp ← sign_extend_if_32bit_op(UserLocal)

        endif
        30: temp ← sign_extend_if_32bit_op(Implementation-Dependent-Value)
        31: temp ← sign_extend_if_32bit_op(Implementation-Dependent-Value)
        otherwise: SignalException(ReservedInstruction)
    endcase
    elseif ((rs==4) and (>=2) and (sel==defined)) // #4 - Release 6
        temp ← sign_extend_if_32bit_op(PerfCtr[sel])
    else
    endif
    GPR[rt] ← temp


```

```
function sign_extend_if_32bit_op(value)
  if (width(value) = 64) and Are64BitOperationsEnabled() then
    sign_extend_if_32bit_op ← value
  else
    sign_extend_if_32bit_op ← sign_extend(value)
  endif
end sign_extend_if_32bit_op
```

Exceptions:

Reserved Instruction

For a register that does not require *sel*, the compiler must support an assembly syntax without *sel* that is ‘RDHWR rt, rs’. Another valid syntax is for *sel* to be 0 to map to pre-Release 6 register numbers which do not require use of *sel* that is, ‘RDHWR rt, rs, 0’.



Format: RDPGPR rt, rs

microMIPS

Purpose: Read GPR from Previous Shadow Set

To move the contents of a GPR from the previous shadow set to a current GPR.

Description: $GPR[rt] \leftarrow SGPR[SRSCtl_{PSS}, rs]$

The contents of the shadow GPR register specified by $SRSCtl_{PSS}$ (signifying the previous shadow set number) and rs (specifying the register number within that set) is moved to the current GPR rt .

Restrictions:

In implementations prior to Release 2 of the Architecture, this instruction resulted in a Reserved Instruction exception.


Operation:

$GPR[rt] \leftarrow SGPR[SRSCtl_{PSS}, rs]$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: RECIP.fmt
 RECIP.S ft, fs
 RECIP.D ft, fs

microMIPS
 microMIPS

Purpose: Reciprocal Approximation

To approximate the reciprocal of an FP value (quickly).

Description: $FPR[ft] \leftarrow 1.0 / FPR[fs]$

The reciprocal of the value in FPR.fs is approximated and placed into FPR.ft. The operand and result are values in format *fmt*.

The numeric accuracy of this operation is implementation dependent. It does not meet the accuracy specified by the IEEE 754 Floating Point standard. The computed result differs from the both the exact result and the IEEE-mandated representation of the exact result by no more than one unit in the least-significant place (ULP).

It is implementation dependent whether the result is affected by the current rounding mode in FCSR.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Availability and Compatibility:

RECIP.S and RECIP.D: Required in all versions of MIPS64 since MIPS64 Release 1. Not available in MIPS32 Release 1. Required in MIPS32 Release 2 and all subsequent versions of MIPS32. When required, required whenever FPU is present, whether a 32-bit or 64-bit FPU, whether in 32-bit or 64-bit FP Register Mode ($FIR_{F64}=0$ or 1, $Status_{FR}=0$ or 1).

Operation:


$StoreFPR(ft, fmt, 1.0 / valueFPR(fs, fmt))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Division-by-zero, Unimplemented Op, Invalid Op, Overflow, Underflow



Format: RINT.fmt
RINT fd, fs

microMIPS32 Release 6

Purpose: Floating-Point Round to Integral

Scalar floating-point round to integral floating point value.

Description: $FPR[fd] \leftarrow \text{round_int}(FPR[fs])$

The scalar floating-point value in the register *fs* is rounded to an integral valued floating-point number in the same format based on the rounding mode bits *RM* in the FPU Control and Status Register *FCSR*. The result is written to *fd*.

The operands and results are values in floating-point data format *fmt*.

The RINT.fmt instruction corresponds to the **roundToIntegralExact** operation in the IEEE Standard for Floating-Point Arithmetic 754TM-2008. The Inexact exception is signaled if the result does not have the same numerical value as the input operand.

The floating point scalar instruction RINT.fmt corresponds to the MSA vector instruction FRINT.df. I.e. RINT.S corresponds to FRINT.W, and RINT.D corresponds to FRINT.D.

Restrictions:

Data-dependent exceptions are possible as specified by the IEEE Standard for Floating-Point Arithmetic 754TM-2008.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.

Operation:

RINT.fmt:

```

if not IsCoprorocessorEnabled(1)
    then SignalException(CoprorocessorUnusable, 1) endif
if not IsFloatingPointImplemented(fmt)
    then SignalException(ReservedInstruction) endif

fin ← ValueFPR(fs,fmt)
ftmp ← RoundIntFP(fin, fmt)
if( fin ≠ ftmp ) SignalFPEException(InExact)
StoreFPR (fd, fmt, ftmp )

function RoundIntFP(tt, n)
    /* Round to integer operation, using rounding mode FCSR.RM*/
endfunction RoundIntFP


```

Exceptions:

Coprorocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Inexact, Overflow, Underflow



Format: ROTR *rt*, *rs*, *sa*

SmartMIPS Crypto, microMIPS

Purpose: Rotate Word Right

To execute a logical right-rotate of a word by a fixed number of bits.

Description: $GPR[rt] \leftarrow GPR[rs] \times(right) \ sa$

The contents of the low-order 32-bit word of GPR *rs* are rotated right; the word result is sign-extended and placed in GPR *rt*. The bit-rotate amount is specified by *sa*.

Restrictions:

If GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rt]) or
  ((ArchitectureRevision() < 2) and (Config3SM = 0)) then
  UNPREDICTABLE
endif
s ← sa
temp ← GPR[rs]s-1..0 || GPR[rs]31..s
GPR[rt] ← sign_extend(temp)

```

Exceptions:

Reserved Instruction



Format: ROTRV rd, rt, rs

SmartMIPS Crypto, microMIPS

Purpose: Rotate Word Right Variable

To execute a logical right-rotate of a word by a variable number of bits.

Description: $GPR[rd] \leftarrow GPR[rt] \times(right) GPR[rs]$

The contents of the low-order 32-bit word of GPR *rt* are rotated right; the word result is sign-extended and placed in GPR *rd*. The bit-rotate amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

If GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rt]) or
  ((ArchitectureRevision() < 2) and (Config3SM = 0)) then
  UNPREDICTABLE
endif
s ← GPR[rs]4..0
temp ← GPR[rt]s-1..0 || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)

```

Exceptions:

Reserved Instruction



Format: ROUND.L.fmt

ROUND.L.S ft, fs

ROUND.L.D ft, fs

microMIPS

microMIPS

Purpose: Floating Point Round to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding to nearest.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and rounded to nearest/even (rounding mode 0). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *ft* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *fd* for long fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model. It is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


$\text{StoreFPR}(ft, L, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation



Format: ROUND.W.fmt
 ROUND.W.S ft, fs
 ROUND.W.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Round to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding to nearest.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format rounding to nearest/even (rounding mode 0). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *ft* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{31}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{31}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{31}-1$
- $-2^{31}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{31}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *fd* for word fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(ft, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation



Format: RSQRT.fmt
 RSQRT.S ft, fs
 RSQRT.D ft, fs

microMIPS
 microMIPS

Purpose: Reciprocal Square Root Approximation

To approximate the reciprocal of the square root of an FP value (quickly).

Description: $FPR[ft] \leftarrow 1.0 / \text{sqrt}(FPR[fs])$

The reciprocal of the positive square root of the value in FPR *fs* is approximated and placed into FPR *ft*. The operand and result are values in format *fmt*.

The numeric accuracy of this operation is implementation dependent; it does not meet the accuracy specified by the IEEE 754 Floating Point standard. The computed result differs from both the exact result and the IEEE-mandated representation of the exact result by no more than two units in the least-significant place (ULP).

The effect of the current *FCSR* rounding mode on the result is implementation dependent.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Availability and Compatibility:

RSQRT.S and RSQRT.D: Required in all versions of MIPS64 since MIPS64 Release 1. Not available in MIPS32 Release 1. Required in MIPS32 Release 2 and all subsequent versions of MIPS32. When required, required whenever FPU is present, whether a 32-bit or 64-bit FPU, whether in 32-bit or 64-bit FP Register Mode ($FIR_{F64}=0$ or 1, $Status_{FR}=0$ or 1).

Operation:


$\text{StoreFPR}(ft, fmt, 1.0 / \text{SquareRoot}(\text{valueFPR}(fs, fmt)))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Division-by-zero, Unimplemented Operation, Invalid Operation, Overflow, Underflow



Format: SB *rt*, *offset*(*base*)

microMIPS

Purpose: Store Byte

To store a byte to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 8-bit byte of GPR *rt* is stored in memory at the location specified by the effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..3 || (pAddr_2..0 xor ReverseEndian3)
bytesel ← vAddr_2..0 xor BigEndianCPU3
datadoubleword ← GPR[rt]_63-8*bytesel..0 || 08*bytesel
StoreMemory(CCA, BYTE, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Bus Error, Address Error, Watch



Format: SBE *rt*, *offset*(*base*)

microMIPS

Purpose: Store Byte EVA

To store a byte to user mode virtual address space when executing in kernel mode.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 8-bit byte of GPR *rt* is stored in memory at the location specified by the effective address. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The SBE instruction functions the same as the SB instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to 1.

Restrictions:

Only usable when access to Coprocessor0 is enabled and when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor ReverseEndian3)
bytesel ← vAddr2..0 xor BigEndianCPU3
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, BYTE, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable,



Format: SC *rt*, *offset* (*base*)

microMIPS

Purpose: Store Conditional Word

To store a word to memory to complete an atomic read-modify-write

Description: if `atomic_update` then `memory[GPR[base] + offset] ← GPR[rt]`, `GPR[rt] ← 1`
else `GPR[rt] ← 0`

The LL and SC instructions provide primitives to implement atomic read-modify-write (RMW) operations on synchronizable memory locations. In Release 5, the behavior of SC is modified when *Config5_{LLB}*=1.

Release 6 (with *Config5_{ULLS}*=1) formalizes support for uncached LL and SC sequences, whereas the pre-Release 6 LL and SC description applies to cached (coherent/non-coherent) memory types. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

The least-significant 32-bit word in GPR *rt* is conditionally stored in memory at the location specified by the aligned effective address. The signed *offset* is added to the contents of GPR *base* to form an effective address.

The SC completes the RMW sequence begun by the preceding LL instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The least-significant 32-bit word of GPR *rt* is stored to memory at the location specified by the aligned effective address.
- A one, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

If either of the following events occurs between the execution of LL and SC, the SC fails:

- A coherent store is completed by another processor or coherent I/O module into the block of synchronizable physical memory containing the word. The size and alignment of the block is implementation-dependent, but it is at least one word and at most the minimum page size.
- A coherent store is executed between an LL and SC sequence on the same processor to the block of synchronizable physical memory containing the word (if *Config5_{LLB}*=1; else whether such a store causes the SC to fail is not predictable).
- An ERET instruction is executed. (Release 5 includes ERETNC, which will not cause the SC to fail.)

Furthermore, an SC must always compare its address against that of the LL. An SC will fail if the aligned address of the SC does not match that of the preceding LL.

A load that executes on the processor executing the LL/SC sequence to the block of synchronizable physical memory containing the word, will not cause the SC to fail (if *Config5_{LLB}*=1; else such a load may cause the SC to fail).

If any of the events listed below occurs between the execution of LL and SC, the SC may fail where it could have succeeded, i.e., success is not predictable. Portable programs should not cause any of these events.

- A load or store executed on the processor executing the LL and SC that is not to the block of synchronizable physical memory containing the word. (The load or store may cause a cache eviction between the LL and SC that results in SC failure. The load or store does not necessarily have to occur between the LL and SC.)

- Any prefetch that is executed on the processor executing the LL and SC sequence (due to a cache eviction between the LL and SC).
- A non-coherent store executed between an LL and SC sequence to the block of synchronizable physical memory containing the word.
- The instructions executed starting with the LL and ending with the SC do not lie in a 2048-byte contiguous region of virtual memory. (The region does not have to be aligned, other than the alignment required for instruction words.)

CACHE operations that are local to the processor executing the LL/SC sequence will result in unpredictable behaviour of the SC if executed between the LL and SC, that is, they may cause the SC to fail where it could have succeeded. Non-local CACHE operations (address-type with coherent CCA) may cause an SC to fail on either the local processor or on the remote processor in multiprocessor or multi-threaded systems. This definition of the effects of CACHE operations is mandated if *Config5_{LLB}*=1. If *Config5_{LLB}*=0, then CACHE effects are implementation-dependent.

The following conditions must be true or the result of the SC is not predictable—the SC may fail or succeed (if *Config5_{LLB}*=1, then either success or failure is mandated, else the result is **UNPREDICTABLE**):

- Execution of SC must have been preceded by execution of an LL instruction.
- An RMW sequence executed without intervening events that would cause the SC to fail must use the same address in the LL and SC. The address is the *same* if the virtual address, physical address, and cacheability & coherency attribute are identical.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the location:

- **Uniprocessor atomicity:** To provide atomic RMW on a single processor, all accesses to the location must be made with memory access type of either *cached noncoherent* or *cached coherent*. All accesses must be to one or the other access type, and they may not be mixed.
- **MP atomicity:** To provide atomic RMW among multiple processors, all accesses to the location must be made with a memory access type of *cached coherent*.
- **I/O System:** To provide atomic RMW with a coherent I/O system, all accesses to the location must be made with a memory access type of *cached coherent*. If the I/O system does not use coherent memory operations, then atomic RMW cannot be provided with respect to the I/O reads and writes.

Release 6 (with *Config5_{ULLS}*=1) formally defines support for uncached LL and SC with the following constraints.

- Both LL and SC must be uncached, and the address must be defined as synchronizable in the system. If the address is non-synchronizable, then this may result in UNPREDICTABLE behavior. The recommended response is that the sub-system report a Bus Error to the processor.
- The use of uncached LL and SC is applicable to any address within the supported address range of the system, or any system configuration, as long as the system implements means to monitor the sequence.
- The SC that ends the sequence may fail locally, but never succeed locally within the processor. When it does not fail locally, the SC must be issued to a “monitor” which is responsible for monitoring the address. This monitor makes the final determination as to whether the SC fails or not, and communicates this to the processor that initiated the sequence.

It is implementation dependent as to what form the monitor takes. It is however differentiated from cached LL and SC which rely on a coherence protocol to make the determination as to whether the sequence succeeds.

- Same processor uncached (but not cached) stores will cause the sequence to fail if the store address matches that of the sequence. A cached store to the same address will cause UNPREDICTABLE behavior.
- Remote cached coherent stores to the same address will cause UNPREDICTABLE behavior.
- Remote cached non-coherent or uncached stores may cause the sequence to fail if they address the external monitor and the monitor makes this determination.

As emphasized above, it is not recommended that software mix memory access types during LL and SC sequences. That is all memory accesses must be of the same type, otherwise this may result in UNPREDICTABLE behavior.

Conditions that cause UNPREDICTABLE behavior for legacy cached LL and SC sequences may also cause such behavior for uncached sequences.

A PAUSE instruction is no-op'd when it is preceded by an uncached LL.

The semantics of an uncached LL/SC atomic operation applies to any uncached CCA including UCA (UnCached Accelerated). An implementation that supports UCA must guarantee that SC does not participate in store gathering and that it ends any gathering initiated by stores preceding the SC in program order when the SC address coincides with a gathering address.

Restrictions:

The addressed location must have a memory access type of *cached noncoherent* or *cached coherent*; if it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config5_{ULS}*=1) extends support to *uncached* types.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
if LLbit then
    StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 063 || LLbit
LLbit ← 0 // if Config5LLB=1, SC always clears LLbit regardless of address match.
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:


LL and SC are used to atomically update memory locations, as shown below.

```
L1:
    LL    T1, (T0) # load counter
    ADDI  T2, T1, 1 # increment
    SC    T2, (T0) # try to store, checking for atomicity
    BEQC  T2, 0, L1 # if not atomic (0), try again
```

Exceptions between the LL and SC cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LL and SC function on a single processor for *cached noncoherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.

As shown in the instruction drawing above, Release 6 implements a 9-bit offset, whereas all release levels lower than Release 6 of the MIPS architecture implement a 16-bit offset.



Format: SCE *rt*, *offset*(*base*)

microMIPS

Purpose: Store Conditional Word EVA

To store a word to user mode virtual memory while operating in kernel mode to complete an atomic read-modify-write.

Description: if `atomic_update` then `memory[GPR[base] + offset] ← GPR[rt]`, `GPR[rt] ← 1` else `GPR[rt] ← 0`

The LL and SC instructions provide primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

Release 6 (with `Config5_ULLS=1`) formalizes support for uncached LLE and SCE sequences. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

The least-significant 32-bit word in GPR *rt* is conditionally stored in memory at the location specified by the aligned effective address. The 9-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The SCE completes the RMW sequence begun by the preceding LLE instruction executed on the processor. To complete the RMW sequence atomically, the following occurs:

- The least-significant 32-bit word of GPR *rt* is stored to memory at the location specified by the aligned effective address.
- A 1, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

If either of the following events occurs between the execution of LL and SC, the SC fails:

- A coherent store is completed by another processor or coherent I/O module into the block of synchronizable physical memory containing the word. The size and alignment of the block is implementation dependent, but it is at least one word and at most the minimum page size.
- An ERET instruction is executed.

If either of the following events occurs between the execution of LLE and SCE, the SCE may succeed or it may fail; the success or failure is not predictable. Portable programs should not cause one of these events.

- A memory access instruction (load, store, or prefetch) is executed on the processor executing the LLE/SCE.
- The instructions executed starting with the LLE and ending with the SCE do not lie in a 2048-byte contiguous region of virtual memory. (The region does not have to be aligned, other than the alignment required for instruction words.)

The following conditions must be true or the result of the SCE is **UNPREDICTABLE**:

- Execution of SCE must have been preceded by execution of an LLE instruction.
- An RMW sequence executed without intervening events that would cause the SCE to fail must use the same address in the LLE and SCE. The address is the same if the virtual address, physical address, and cacheability & coherency attribute are identical.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LLE/SCE semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the

location:

- **Uniprocessor atomicity:** To provide atomic RMW on a single processor, all accesses to the location must be made with memory access type of either *cached non coherent* or *cached coherent*. All accesses must be to one or the other access type, and they may not be mixed.
- **MP atomicity:** To provide atomic RMW among multiple processors, all accesses to the location must be made with a memory access type of *cached coherent*.
- **I/O System:** To provide atomic RMW with a coherent I/O system, all accesses to the location must be made with a memory access type of *cached coherent*. If the I/O system does not use coherent memory operations, then atomic RMW cannot be provided with respect to the I/O reads and writes.

The SCE instruction functions the same as the SC instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to 1.

The definition for SCE is extended for uncached memory types in a manner identical to SC. The extension is defined in the SC instruction description.

Restrictions:

The addressed location must have a memory access type of *cached non coherent* or *cached coherent*; if it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config5_{ULS}*=1) extends support to *uncached* types.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
if LLbit then
    StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 063 || LLbit

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch, Reserved Instruction, Coprocessor Unusable

Programming Notes:

LLE and SCE are used to atomically update memory locations, as shown below.


```

L1:
    LLE    T1, (T0)  # load counter
    ADDI   T2, T1, 1 # increment
    SCE    T2, (T0)  # try to store, checking for atomicity
    BEQC   T2, 0, L1 # if not atomic (0), try again

```


Exceptions between the LLE and SCE cause SCE to fail, so persistent exceptions must be avoided. Examples are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LLE and SCE function on a single processor for *cached non coherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.



Format: SCD *rt*, *offset*(*base*)

microMIPS

Purpose: Store Conditional Doubleword

To store a doubleword to memory to complete an atomic read-modify-write.

Description: if *atomic_update* then $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}], \text{GPR}[\text{rt}] \leftarrow 1$ else $\text{GPR}[\text{rt}] \leftarrow 0$

The LLD and SCD instructions provide primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

Release 6 (with *Config5_{ULS}* =1) formalizes support for uncached LLD and SCD sequences, whereas the pre-Release 6 LLD and SCD description applies to cached (coherent/non-coherent) memory types. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

The 64-bit doubleword in GPR *rt* is conditionally stored in memory at the location specified by the aligned effective address. The signed *offset* is added to the contents of GPR *base* to form an effective address.

The SCD completes the RMW sequence begun by the preceding LLD instruction executed on the processor. If SCD completes the RMW sequence atomically, the following occurs:

- The 64-bit doubleword of GPR *rt* is stored into memory at the location specified by the aligned effective address.
- A 1, indicates success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

If either of the following events occurs between the execution of LLD and SCD, the SCD fails:

- A coherent store is completed by another processor or coherent I/O module into the block of synchronizable physical memory containing the doubleword. The size and alignment of the block is implementation dependent, but it is at least one doubleword and at most the minimum page size.
- An ERET instruction is executed.

If either of the following events occurs between the execution of LLD and SCD, the SCD may succeed or it may fail; the success or failure is not predictable. Portable programs should not cause the following events:

- A memory access instruction (load, store, or prefetch) is executed on the processor executing the LLD/SCD.
- The instructions executed starting with the LLD and ending with the SCD do not lie in a 2048-byte contiguous region of virtual memory. (The region does not have to be aligned, other than the alignment required for instruction words.)

The following two conditions must be true or the result of the SCD is **UNPREDICTABLE**:

- Execution of the SCD must be preceded by execution of an LLD instruction.

- An RMW sequence executed without intervening events that would cause the SCD to fail must use the same address in the LLD and SCD. The address is the same if the virtual address, physical address, and cache-coherence algorithm are identical.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the location:

- **Uniprocessor atomicity:** To provide atomic RMW on a single processor, all accesses to the location must be made with memory access type of either *cached non coherent* or *cached coherent*. All accesses must be to one or the other access type, and they may not be mixed.
- **MP atomicity:** To provide atomic RMW among multiple processors, all accesses to the location must be made with a memory access type of *cached coherent*.
- **I/O System:** To provide atomic RMW with a coherent I/O system, all accesses to the location must be made with a memory access type of *cached coherent*. If the I/O system does not use coherent memory operations, then atomic RMW cannot be provided with respect to the I/O reads and writes.

Release 6 (with *Config_{ULS}*=1) formally defines support for uncached LLD and SCD with the following constraints.

- Both LLD and SCD must be uncached, and the address must be defined as synchronizable in the system. If the address is non-synchronizable, then this may result in UNPREDICTABLE behavior. The recommended response is that the sub-system report a Bus Error to the processor.
- The use of uncached LLD and SCD is applicable to any address within the supported address range of the system, or any system configuration, as long as the system implements means to monitor the sequence.
- The SCD that ends the sequence may fail locally, but never succeed locally within the processor. When it does not fail locally, the SCD must be issued to a “monitor” which is responsible for monitoring the address. This monitor makes the final determination as to whether the SCD fails or not, and communicates this to the processor that initiated the sequence.

It is implementation dependent as to what form the monitor takes. It is however differentiated from cached LLD and SCD which rely on a coherence protocol to make the determination as to whether the sequence succeeds.

- Same processor uncached (but not cached) stores will cause the sequence to fail if the store address matches that of the sequence. A cached store to the same address will cause UNPREDICTABLE behavior.
- Remote cached coherent stores to the same address will cause UNPREDICTABLE behavior.
- Remote cached non-coherent or uncached stores may cause the sequence to fail if they address the external monitor and the monitor makes this determination.

As emphasized above, it is not recommended that software mix memory access types during LLD and SCD sequences. That is all memory accesses must be of the same type, otherwise this may result in UNPREDICTABLE behavior.

Conditions that cause UNPREDICTABLE behavior for legacy cached LLD and SCD sequences may also cause such behavior for uncached sequences.

A PAUSE instruction is no-op'd when it is preceded by an uncached LLD.

The semantics of an uncached LLD/SCD atomic operation applies to any uncached CCA including UCA (UnCached Accelerated). An implementation that supports UCA must guarantee that SCD does not participate in store gathering and that it ends any gathering initiated by stores preceding the SCD in program order when the SCD address coincides with a gathering address.

Restrictions:

The addressed location must have a memory access type of *cached non coherent* or *cached coherent*; if it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config_{ULS}*=1) extends support to *uncached* types.

The effective address must be naturally-aligned. If any of the 3 least-significant bits of the address is non-zero, an Address Error exception occurs.

Providing misaligned support for Release 6 is not a requirement for this instruction.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
datadoubleword ← GPR[rt]
if LLbit then
    StoreMemory (CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 063 || LLbit

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Reserved Instruction, Watch

Programming Notes:

LLD and SCD are used to atomically update memory locations, as shown below.

```

L1:
    LLD    T1, (T0)  # load counter
    ADDI   T2, T1, 1 # increment
    SCD    T2, (T0)  # try to store,
                    # checking for atomicity
    BEQC   T2, 0, L1 # if not atomic (0), try again

```

Exceptions between the LLD and SCD cause SCD to fail, so persistent exceptions must be avoided. Examples are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LLD and SCD function on a single processor for *cached non coherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.

31	26	25	21	20	16	15	12	11	9	8	4	3	0
POOL32C 011000	rt		base		SCDP 1101		0 000		rd		0 0000		
6	5		5		4		3		5		4		

Format: SCDP rt, rd, (base)

microMIPS64 Release 6

Purpose: Store Conditional DoubleWord Paired

Conditionally store a paired double-word to memory to complete an atomic read-modify-write

Description: if atomic_update then memory[GPR[base]] ← {GPR[rd], GPR[rt]}, GPR[rt] ← 1
else GPR[rt] ← 0

The LLDP and SCDP instructions provide primitives to implement a paired double-word atomic read-modify-write (RMW) operation at a synchronizable memory location.

Release 6 (with *Config_{ULLS}*=1) formalizes support for uncached LLDP and SCDP sequences. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

A paired double-word is conditionally written to memory in a single atomic memory operation. GPR *rd* is the most-significant double-word and GPR *rt* is the least-significant double-word of the quad-word in memory. The write occurs to a quad-word aligned effective address from GPR *base*.

A paired double-word read or write occurs as a pair of double-word reads or writes that is quad-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The SCDP completes the RMW sequence begun by the preceding LLDP instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The paired double-word formed from the concatenation of GPRs *rd* and *rt* is stored to memory at the location specified by the quad-word aligned effective address.
- A one, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

Though legal programming requires LLDP to start the atomic read-modify-write sequence and SCDP to end the same sequence, whether the SCDP completes is only dependent on the state of LLbit and *LLAddr*, which are set by a preceding load-linked instruction of any type. Software must assume that pairing load-linked and store-conditional instructions in an inconsistent manner causes **UNPREDICTABLE** behavior.

The SCDP must always compare its quad-word aligned address against that of the preceding LLDP. The SCDP will fail if the address does not match that of the preceding LLDP.

Events that occur between the execution of load-linked and store-conditional instruction types that must cause the sequence to fail are given in the legacy SCD instruction definition, *except the block of synchronizable memory is a quadword, not doubleword*.

Additional events that occur between the execution of load-linked and store-conditional instruction types that *may* cause success of the sequence to be **UNPREDICTABLE** are defined in the SCD instruction definition.

A load that executes on the processor executing the LLDP/SCDP sequence to the block of synchronizable physical memory containing the paired double-word, will not cause the SCDP to fail.

Effect of CACHE operations, both local and remote, on a paired double-word atomic operation are defined in the SC instruction definition.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the location. Requirements for Uniprocessor, MP and I/O atomicity are given in the SC definition.

The definition for SCDP is extended for uncached memory types in a manner identical to SC. The extension is defined in the SC instruction description.

Restrictions:

Load-Linked and Store-Conditional instruction types *require* that the addressed location must have a memory access type of *cached noncoherent* or *cached coherent*, that is the processor must have a cache. If it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config5_{ULS}*=1) extends support to *uncached* types.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if *Config5_{XNP}*=0.

Operation:

```
vAddr ← GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
dataquadword ← {GPR[rd], GPR[rt]}
if (LLbit && (pAddr == LLAddr)) then // quadword aligned monitor
    // PAIREDDOUBLEWORD: two double-word data-type that is quad-word atomic
    StoreMemory (CCA, PAIREDDOUBLEWORD, dataquadword, pAddr, vAddr, DATA)
    GPR[rt] ← 063 || 1'b1
else
    GPR[rt] ← 064
endif
LLbit ← 0
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Reserved Instruction, Address Error, Watch

Programming Notes:


LLDP and SCDP are used to atomically update memory locations, as shown below.

```
L1:
    LLDP    T2, T3, (T0) # load T2 and T3
    BOVC    T2, 1, U32 # check whether least-significant double-word may overflow
    ADDI     T2, T2, 1 # increment lower - only
    SCDP     T2, T3, (T0) # store T2 and T3
    BEQC     T2, 0, L1 # if not atomic (0), try again

U32:
    ADDI     T2, T2, 1 # increment lower
    ADDI     T3, T3, 1 # increment upper
    SCDP     T2, T3, (T0)
    BEQC     T2, 0, L1 # if not atomic (0), try again
```

Exceptions between the LLDP and SCDP cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LLDP and SCDP function on a single processor for *cached noncoherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.



Format: SCWP *rt*, *rd*, (*base*)

microMIPS Release 6

Purpose: Store Conditional Word Paired

Conditionally store a paired word to memory to complete an atomic read-modify-write.

Description: if `atomic_update` then `memory[GPR[base]] ← {GPR[rd], GPR[rt]}`, `GPR[rt] ← 1`
else `GPR[rt] ← 0`

The LLWP and SCWP instructions provide primitives to implement a paired word atomic read-modify-write (RMW) operation at a synchronizable memory location.

Release 6 (with *Config5_{ULS}* = 1) formalizes support for uncached LLWP and SCWP sequences. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

A paired word is formed from the concatenation of GPR *rd* and GPR *rt*. GPR *rd* is the most-significant word of the paired word, and GPR *rt* is the least-significant word of the paired word. The paired word is conditionally stored in memory at the location specified by the double-word aligned effective address from GPR *base*.

A paired word read or write occurs as a pair of word reads or writes that is double-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The SCWP completes the RMW sequence begun by the preceding LLWP instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The paired word formed from the concatenation of GPRs *rd* and *rt* is stored to memory at the location specified by the double-word aligned effective address.
- A one, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

Though legal programming requires LLWP to start the atomic read-modify-write sequence and SCWP to end the same sequence, whether the SCWP completes is only dependent on the state of LLbit and *LLAddr*, which are set by a preceding load-linked instruction of any type. Software must assume that pairing load-linked and store-conditional instructions in an inconsistent manner causes **UNPREDICTABLE** behavior.

The SCWP must always compare its double-word aligned address against that of the preceding LLWP. The SCWP will fail if the address does not match that of the preceding LLWP.

Events that occur between the execution of load-linked and store-conditional instruction types that *must* cause the sequence to fail are given in the legacy SC instruction description.

Additional events that occur between the execution of load-linked and store-conditional instruction types that *may* cause success of the sequence to be **UNPREDICTABLE** are defined in the SC instruction description.

A load that executes on the processor executing the LLWP/SCWP sequence to the block of synchronizable physical memory containing the paired word, will not cause the SCWP to fail.

Effect of CACHE operations, both local and remote, on a paired word atomic operation are defined in the SC instruction description.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the location. Requirements for Uniprocessor, MP and I/O atomicity are given in the SC definition.

The definition for SCWP is extended for uncached memory types in a manner identical to SC. The extension is defined in the SC instruction description.

Restrictions:

Load-Linked and Store-Conditional instruction types *require* that the addressed location must have a memory access type of *cached noncoherent* or *cached coherent*, that is the processor must have a cache. If it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config5_{ULS}*=1) extends support to *uncached* types.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if *Config5_{XNP}*=0.

Operation:

```
vAddr ← GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
datadoubleword31..0 ← GPR[rt]31..0
datadoubleword63..32 ← GPR[rd]31..0
if (LLbit && (pAddr == LLAddr))then
  // PAIREDWORD: two word data-type that is double-word atomic
  StoreMemory (CCA, PAIREDWORD, datadoubleword, pAddr, vAddr, DATA)
  GPR[rt] ← 063 || 1'b1
else
  GPR[rt] ← 064
endif
LLbit ← 0
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Reserved Instruction, Address Error, Watch

Programming Notes:

LLWP and SCWP are used to atomically update memory locations, as shown below.

```
L1:
  LLWP  T2, T3, (T0)  # load T2 and T3
  BOVC  T2, 1, U32    # check whether least-significant word may overflow
  ADDI   T2, T2, 1     # increment lower - only
  SCWP   T2, T3, (T0)  # store T2 and T3
  BEQC   T2, 0, L1     # if not atomic (0), try again

U32:
  ADDI   T2, T2, 1     # increment lower
  ADDI   T3, T3, 1     # increment upper
  SCWP   T2, T3, (T0)
  BEQC   T2, 0, L1     # if not atomic (0), try again
```

Exceptions between the LLWP and SCWP cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LLWP and SCWP function on a single processor for *cached noncoherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.

31	26	25	21	20	16	15	12	11	9	8	4	3	0				
POOL32C 011000						rt		base		ST-EVA 1010		SCWPE 000		rd		0 0000	
6						5		5		4		3		5		4	

Format: SCWPE *rt*, *rd*, (*base*)

microMIPS Release 6

Purpose: Store Conditional Word Paired EVA

Conditionally store a paired word to memory to complete an atomic read-modify-write. The store occurs in kernel mode to user virtual address space.

Description: if `atomic_update` then `memory[GPR[base]] ← {GPR[rd], GPR[rt]}`, `GPR[rt] ← 1`
else `GPR[rt] ← 0`

The LLWPE and SCWPE instructions provide primitives to implement a paired word atomic read-modify-write (RMW) operation at a synchronizable memory location.

Release 6 (with *Config5_{ULS}*=1) formalizes support for uncached LLWPE and SCWPE sequences. (The description for uncached support does not modify the description for cached support and is written in a self-contained manner.)

A paired word is formed from the concatenation of GPR *rd* and GPR *rt*. GPR *rd* is the most-significant word of the double-word, and GPR *rt* is the least-significant word of the double-word. The paired word is conditionally stored in memory at the location specified by the double-word aligned effective address from GPR *base*.

A paired word read or write occurs as a pair of word reads or writes that is double-word atomic.

The instruction has no offset. The effective address is equal to the contents of GPR *base*.

rd is intentionally positioned in a non-standard bit-range.

The SCWPE completes the RMW sequence begun by the preceding LLWPE instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The paired word formed from the concatenation of GPRs *rd* and *rt* is stored to memory at the location specified by the double-word aligned effective address.
- A one, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

Though legal programming requires LLWPE to start the atomic read-modify-write sequence and SCWPE to end the same sequence, whether the SCWPE completes is only dependent on the state of LLbit and *LLAddr*, which are set by a preceding load-linked instruction of any type. Software must assume that pairing load-linked and store-conditional instructions in an inconsistent manner causes **UNPREDICTABLE** behavior.

The SCWPE must always compare its double-word aligned address against that of the preceding LLWPE. The SCWPE will fail if the address does not match that of the preceding LLWPE.

The SCWPE instruction functions the same as the SCWP instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Segmentation Control for additional information.

Events that occur between the execution of load-linked and store-conditional instruction types that must cause the sequence to fail are given in the legacy SC instruction definition..

Additional events that occur between the execution of load-linked and store-conditional instruction types that *may* cause success of the sequence to be **UNPREDICTABLE** are defined in the SC instruction definition.

A load that executes on the processor executing the LLWPE/SCWPE sequence to the block of synchronizable physical memory containing the paired word, will not cause the SCWPE to fail.

Effect of CACHE operations, both local and remote, on a paired word atomic operation are defined in the SC instruction definition.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the location. Requirements for Uniprocessor, MP and I/O atomicity are given in the SC definition.

The definition for SCWPE is extended for uncached memory types in a manner identical to SC. The extension is defined in the SC instruction description.

Restrictions:

Load-Linked and Store-Conditional instruction types *require* that the addressed location must have a memory access type of *cached noncoherent* or *cached coherent*, that is the processor must have a cache. If it does not, the result is **UNPREDICTABLE**. Release 6 (with *Config5_{ULLS}*=1) extends support to *uncached* types.

The architecture optionally allows support for Load-Linked and Store-Conditional instruction types in a cacheless processor. Support for cacheless operation is implementation dependent. In this case, *LLAddr* is optional.

Providing misaligned support is not a requirement for this instruction.

Availability and Compatibility

This instruction is introduced by Release 6. It is only present if *Config5_{XNP}*=0 and *Config5_{EVA}*=1.

Operation:

```
vAddr ← GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
datadoubleword31..0 ← GPR[rt]31..0
datadoubleword63..32 ← GPR[rd]31..0
if (LLbit && (pAddr == LLAddr))then
    // PAIREDWORD: two word data-type that is double-word atomic
    StoreMemory (CCA, PAIREDWORD, datadoubleword, pAddr, vAddr, DATA)
    GPR[rt] ← 063 || 1'b1
else
    GPR[rt] ← 064
endif
LLbit ← 0
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Reserved Instruction, Address Error, Watch, Coprocessor Unusable.

Programming Notes:


LLWPE and SCWPE are used to atomically update memory locations, as shown below.

```
L1:
    LLWPE  T2, T3, (T0)    # load T2 and T3
    BOVC   T2, 1, U32      # check whether least-significant word may overflow
    ADDI    T2, T2, 1       # increment lower - only
    SCWPE  T2, T3, (T0)    # store T2 and T3
    BEQC    T2, 0, L1      # if not atomic (0), try again
```

```
U32:
    ADDI    T2, T2, 1      # increment lower
    ADDI    T3, T3, 1      # increment upper
    SCWPE   T2, T3, (T0)
    BEQC    T2, 0, L1      # if not atomic (0), try again
```

Exceptions between the LLWPE and SCWPE cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LLWPE and SCWPE function on a single processor for *cached noncoherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.



Format: SD *rt*, *offset*(*base*)

microMIPS64

Purpose: Store Doubleword

To store a doubleword to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The 64-bit doubleword in GPR *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-release 6: The effective address must be naturally-aligned. If any of the 3 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
datadoubleword ← GPR[rt]
StoreMemory (CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Reserved Instruction, Watch

31	26	25	16	15	6	5	0
POOL32A 000000	code - use syscall				SDBBP 1101101101	POOL32AXf 111100	
6	10				10	6	

Format: SDBBP code

EJTAG microMIPS

Purpose: Software Debug Breakpoint

To cause a debug breakpoint exception

Description:

This instruction causes a debug exception, passing control to the debug exception handler. If the processor is executing in Debug Mode when the SDBBP instruction is executed, the exception is a Debug Mode Exception, which sets the Debug_{DExcCode} field to the value 0x9 (Bp). The code field can be used for passing information to the debug exception handler, and is retrieved by the debug exception handler only by loading the contents of the memory word containing the instruction, using the DEPC register. The CODE field is not used in any way by the hardware.


Restrictions:

Operation:

```
if Config5.SBRI=1 then /* SBRI is a MIPS Release 6 feature */
    SignalException(ReservedInstruction) endif
If DebugDM = 1 then SignalDebugModeBreakpointException() endif // nested
SignalDebugBreakpointException() // normal
```

Exceptions:

Debug Breakpoint Exception
Debug Mode Breakpoint Exception



Format: SDC1 ft, offset(base)

microMIPS

Purpose: Store Doubleword from Floating Point

To store a doubleword from an FPR to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{FPR}[\text{ft}]$

The 64-bit doubleword in FPR *ft* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
datadoubleword ← ValueFPR(ft, UNINTERPRETED_DOUBLEWORD)
StoreMemory(CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SDC2 rt, offset(base)

microMIPS

Purpose: Store Doubleword from Coprocessor 2

To store a doubleword from a Coprocessor 2 register to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{CPR}[2, \text{rt}, 0]$

The 64-bit doubleword in Coprocessor 2 register *rt* is stored in memory at the location specified by the aligned effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
datadoubleword ← CPR[2,rt,0]
StoreMemory(CCA, DOUBLEWORD, datadoubleword, pAddr, vAddr, DATA)


```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:

As shown in the instruction drawing above, Release 6 implements an 11-bit offset, whereas all release levels lower than Release 6 of the MIPS architecture implement a 16-bit offset.



Format: SEB *rt*, *rs*

microMIPS

Purpose: Sign-Extend Byte

To sign-extend the least significant byte of GPR *rs* and store the value into GPR *rt*.

Description: $\text{GPR}[rt] \leftarrow \text{SignExtend}(\text{GPR}[rs]_{7..0})$

The least significant byte from GPR *rs* is sign-extended and stored in GPR *rt*.

Restrictions:

Prior to architecture Release 2, this instruction resulted in a Reserved Instruction exception.

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
GPR[rt] ← sign_extend(GPR[rs]_{7..0})

```

Exceptions:


Reserved Instruction

Programming Notes:

For symmetry with the SEB and SEH instructions, you expect that there would be ZEB and ZEH instructions that zero-extend the source operand and expect that the SEW and ZEW instructions would exist to sign- or zero-extend a word to a doubleword. These instructions do not exist because there are functionally-equivalent instructions already in the instruction set. The following table shows the instructions providing the equivalent functions.

Expected Instruction	Function	Equivalent Instruction
ZEB <i>rx</i> , <i>ry</i>	Zero-Extend Byte	ANDI <i>rx</i> , <i>ry</i> , 0xFF
ZEH <i>rx</i> , <i>ry</i>	Zero-Extend Halfword	ANDI <i>rx</i> , <i>ry</i> , 0xFFFF
SEW <i>rx</i> , <i>ry</i>	Sign-Extend Word	SLL <i>rx</i> , <i>ry</i> , 0
ZEW <i>rx</i> , <i>rx</i> ¹	Zero-Extend Word	DINSP32 <i>rx</i> , <i>r0</i> , 32, 32

1. The equivalent instruction uses *rx* for both source and destination, so the expected instruction is limited to one register



Format: SEH *rt*, *rs*

microMIPS

Purpose: Sign-Extend Halfword

To sign-extend the least significant halfword of GPR *rs* and store the value into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{SignExtend}(GPR[rs]_{15..0})$

The least significant halfword from GPR *rs* is sign-extended and stored in GPR *rt*.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
GPR[rt] ← sign_extend(GPR[rs]_{15..0})

```

Exceptions:

Reserved Instruction

Programming Notes:

The SEH instruction can be used to convert two contiguous halfwords to sign-extended word values in three instructions. For example:

```

lw    t0, 0(a1)           /* Read two contiguous halfwords */
seh    t1, t0              /* t1 = lower halfword sign-extended to word */
sra    t0, t0, 16          /* t0 = upper halfword sign-extended to word */


```

Zero-extended halfwords can be created by changing the SEH and SRA instructions to ANDI and SRL instructions, respectively.

For symmetry with the SEB and SEH instructions, you expect that there would be ZEB and ZEH instructions that zero-extend the source operand and expect that the SEW and ZEW instructions would exist to sign- or zero-extend a word to a doubleword. These instructions do not exist because there are functionally-equivalent instructions already in the instruction set. The following table shows the instructions providing the equivalent functions.

Expected Instruction	Function	Equivalent Instruction
ZEB <i>rx</i> , <i>ry</i>	Zero-Extend Byte	ANDI <i>rx</i> , <i>ry</i> , 0xFF
ZEH <i>rx</i> , <i>ry</i>	Zero-Extend Halfword	ANDI <i>rx</i> , <i>ry</i> , 0xFFFF
SEW <i>rx</i> , <i>ry</i>	Sign-Extend Word	SLL <i>rx</i> , <i>ry</i> , 0
ZEW <i>rx</i> , <i>rx</i> ¹	Zero-Extend Word	DINSP32 <i>rx</i> , <i>r0</i> , 32, 32

1. The equivalent instruction uses rx for both source and destination, so the expected instruction is limited to one register



Format: SEL.fmt
SEL fd, fs, ft, fmt

microMIPS32 Release 6

Purpose: Select floating point values with FPR condition

Description: $FPR[fd] \leftarrow FPR[fd].bit0 ? FPR[ft] : FPR[fs]$

SEL.fmt is a select operation, with a condition input in FPR fd, and 2 data inputs in FPRs ft and fs.

- If the condition is true, the value of ft is written to fd.
- If the condition is false, the value of fs is written to fd.

The condition input is specified by FPR fd, and is overwritten by the result.

The condition is true only if bit 0 of the condition input FPR fd is set. Other bits are ignored.

This instruction has floating point formats S and D, but these specify only the width of the operands. SEL.S can be used for 32-bit W data, and SEL.D can be used for 64 bit L data.

This instruction does not cause data-dependent exceptions. It does not trap on NaNs. It does not set the FPU Cause bits.

Restrictions:

None

Availability and Compatibility:

SEL.fmt is introduced by and required as of microMIPS32 Release 6.

Special Considerations:

Only formats S and D are valid. Other format values may be used to encode other instructions. Unused format encodings are required to signal the Reserved Instruction exception.

Operation:

```
tmp ← ValueFPR(fd, UNINTERPRETED_WORD)
cond ← tmp.bit0
if cond then
    tmp ← ValueFPR(ft, fmt)
else
    tmp ← ValueFPR(fs, fmt)
endif
StoreFPR(fd, fmt, tmp)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

None

31	26	25	21	20	16	15	11	10	6	5	0
POOL32A 000000	rt		rs		rd		0		SELEQZ 0101000000		
POOL32A 000000	rt		rs		rd		0		SELNEZ 0110000000		
6	5		5		5		1		10		

Format: SELEQZ SELNEZ
 SELEQZ rd,rs,rt
 SELNEZ rd,rs,rt

microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Select integer GPR value or zero

Description:

SELEQZ: $GPR[rd] \leftarrow GPR[rt] ? 0 : GPR[rs]$
 SELNEZ: $GPR[rd] \leftarrow GPR[rt] ? GPR[rs] : 0$

- SELEQZ is a select operation, with a condition input in GPR *rt*, one explicit data input in GPR *rs*, and implicit data input 0. The condition is true only if all bits in GPR *rt* are zero.
- SELNEZ is a select operation, with a condition input in GPR *rt*, one explicit data input in GPR *rs*, and implicit data input 0. The condition is true only if any bit in GPR *rt* is nonzero

If the condition is true, the value of *rs* is written to *rd*.

If the condition is false, the zero is written to *rd*.

This instruction operates on all GPRLEN bits of the CPU registers, that is, all 32 bits on a 32-bit CPU, and all 64 bits on a 64-bit CPU. All GPRLEN bits of *rt* are tested.

Restrictions:

None

Availability and Compatibility:

These instructions are introduced by and required as of MIPS32 Release 6.

Special Considerations:

None

Operation:

```
SELNEZ: cond ← GPR[rt] ≠ 0
SELEQZ: cond ← GPR[rt] = 0
if cond then
    tmp ← GPR[rs]
else
    tmp ← 0
endif
GPR[rd] ← tmp
```

Exceptions:

None

Programming Note:

Release 6 removes the Pre-Release 6 instructions MOVZ and MOVN:

```
MOVZ: if GPR[rt] = 0 then GPR[rd] ← GPR[rs]
MOVN: if GPR[rt] ≠ 0 then GPR[rd] ← GPR[rs]
```

MOVZ can be emulated using Release 6 instructions as follows:

```
SELEQZ at, rs, rt
SELNEZ rd, rd, rt
OR rd, rd, at
```

Similarly MOVN:

```
SELNEZ at, rs, rt
SELEQZ rd, rd, rt
OR rd, rd, at
```

The more general select operation requires 4 registers (1 output + 3 inputs (1 condition + 2 data)) and can be expressed:

```
rD ← if rC then rA else rB
```

The more general select can be created using Release 6 instructions as follows:

```
SELNEZ at, rB, rC
SELNEZ rD, rA, rC
OR rD, rD, at
```


31	26	25	21	20	16	15	11	10	6	5	0
POOL32F 010101	ft		fs		fd		fmt		SELEQZ.fmt 000111000		
POOL32F 010101	ft		fs		fd		fmt		SELNEZ.fmt 001111000		
6	5		5		5		2		9		

Format: SELEQZ.fmt SELNEQZ.fmt
 SELEQZ.S fd,fs,ft
 SELEQZ.D fd,fs,ft
 SELNEZ.S fd,fs,ft
 SELNEZ.D fd,fs,ft

microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6
 microMIPS32 Release 6

Purpose: Select floating point value or zero with FPR condition.

Description:

```
SELEQZ.fmt: FPR[fd] ← FPR[ft].bit0 ? 0 : FPR[fs]
SELNEZ.fmt: FPR[fd] ← FPR[ft].bit0 ? FPR[fs] : 0
```

- SELEQZ.fmt is a select operation, with a condition input in FPR ft, one explicit data input in FPR fs, and implicit data input 0. The condition is true only if bit 0 of FPR ft is zero.
- SELNEZ.fmt is a select operation, with a condition input in FPR ft, one explicit data input in FPR fs, and implicit data input 0. The condition is true only if bit 0 of FPR ft is nonzero.

If the condition is true, the value of fs is written to fd.

If the condition is false, the value that has all bits zero is written to fd.

This instruction has floating point formats S and D, but these specify only the width of the operands. Format S can be used for 32-bit W data, and format D can be used for 64 bit L data. The condition test is restricted to bit 0 of FPR ft. Other bits are ignored.

This instruction has no execution exception behavior. It does not trap on NaNs. It does not set the FPU Cause bits.

Restrictions:

FPR fd destination register bits beyond the format width are UNPREDICTABLE. For example, if fmt is S, then fd bits 0-31 are defined, but bits 32 and above are UNPREDICTABLE. If fmt is D, then fd bits 0-63 are defined.

Availability and Compatibility:

These instructions are introduced by and required as of MIPS32 Release 6.

Special Considerations:

Only formats S and D are valid. Other format values may be used to encode other instructions. Unused format encodings are required to signal the Reserved Instruction exception.

Operation:


```
tmp ← ValueFPR(ft, UNINTERPRETED_WORD)
SELEQZ: cond ← tmp.bit0 = 0
SELNEZ: cond ← tmp.bit0 ≠ 0
if cond then
    tmp ← ValueFPR(fs, fmt)
else
    tmp ← 0 /* all bits set to zero */
```

```
endif  
StoreFPR(fd, fmt, tmp)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:



Format: SH *rt*, *offset* (*base*)

microMIPS

Purpose: Store Halfword

To store a halfword to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 16-bit halfword of register *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian2 || 0))
bytesel ← vAddr2..0 xor (BigEndianCPU2 || 0)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, HALFWORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SHE rt, offset(base)

microMIPS

Purpose: Store Halfword EVA

To store a halfword to user mode virtual address space when executing in kernel mode.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 16-bit halfword of register *rt* is stored in memory at the location specified by the aligned effective address. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The SHE instruction functions the same as the SH instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to 1.

Restrictions:

Only usable in kernel mode when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Pre-Release 6: The effective address must be naturally-aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian2 || 0))
bytesel ← vAddr2..0 xor (BigEndianCPU2 || 0)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, HALFWORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



Format: SIGRIE code

MIPS32 Release 6

Purpose: Signal Reserved Instruction Exception

The SIGRIE instruction signals a Reserved Instruction exception.

Description: `SignalException(ReservedInstruction)`

The SIGRIE instruction signals a Reserved Instruction exception. Implementations should use exactly the same mechanisms as they use for reserved instructions that are not defined by the Architecture.

The 16-bit *code* field is available for software use.

Restrictions:

The 16-bit *code* field is available for software use. The value zero is considered the default value. Software may provide extended functionality by interpreting nonzero values of the *code* field in a manner that is outside the scope of this architecture specification.

Availability and Compatibility:

This instruction is introduced by and required as of Release 6.


Pre-Release 6: this instruction encoding was reserved, and required to signal a Reserved Instruction exception. Therefore this instruction can be considered to be both backwards and forwards compatible.

Operation:

`SignalException(ReservedInstruction)`

Exceptions:

Reserved Instruction



Format: SLL *rt*, *rs*, *sa*

microMIPS

Purpose: Shift Word Left Logical

To left-shift a word by a fixed number of bits.

Description: $GPR[rt] \leftarrow GPR[rs] \ll sa$

The contents of the low-order 32-bit word of GPR *rs* are shifted left, inserting zeros into the emptied bits. The word result is sign-extended and placed in GPR *rt*. The bit-shift amount is specified by *sa*.

Restrictions:

None

Operation:

```
s ← sa
temp ← GPR[rs] (31-s) .. 0 || 0s
GPR[rt] ← sign_extend(temp)
```

Exceptions:


None

Programming Notes:

The SLL input operand does not have to be a properly sign-extended word value to produce a valid sign-extended 32-bit result. The result word is always sign-extended into a 64-bit destination register; this instruction with a zero shift amount truncates a 64-bit value to 32 bits and sign-extends it.

SLL *r0*, *r0*, 0, expressed as NOP, is the assembly idiom used to denote no operation.

SLL *r0*, *r0*, 1, expressed as SSNOP, is the assembly idiom used to denote no operation that causes an issue break on superscalar processors.



Format: SLLV rd, rt, rs

microMIPS

Purpose: Shift Word Left Logical Variable

To left-shift a word by a variable number of bits.

Description: $GPR[rd] \leftarrow GPR[rt] \ll GPR[rs]$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeros into the emptied bits. The resulting word is sign-extended and placed in GPR *rd*. The bit-shift amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

None

Operation:

```

s ← GPR[rs]4..0
temp ← GPR[rt](31-s)..0 || 0s
GPR[rd] ← sign_extend(temp)


```

Exceptions:

None

Programming Notes:

The input operand does not have to be a properly sign-extended word value to produce a valid sign-extended 32-bit result. The result word is always sign-extended into a 64-bit destination register; this instruction with a zero shift amount truncates a 64-bit value to 32 bits and sign-extends it.



Format: SLT rd, rs, rt

microMIPS

Purpose: Set on Less Than

To record the result of a less-than comparison.

Description: $GPR[rd] \leftarrow (GPR[rs] < GPR[rt])$

Compare the contents of GPR *rs* and GPR *rt* as signed integers; record the Boolean result of the comparison in GPR *rd*. If GPR *rs* is less than GPR *rt*, the result is 1 (true); otherwise, it is 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None

Operation:


```

if GPR[rs] < GPR[rt] then
    GPR[rd] ← 0GPRLEN-1 || 1
else
    GPR[rd] ← 0GPRLEN
endif

```

Exceptions:

None



Format: SLTI *rt*, *rs*, *immediate*

microMIPS

Purpose: Set on Less Than Immediate

To record the result of a less-than comparison with a constant.

Description: $GPR[rt] \leftarrow (GPR[rs] < \text{sign_extend}(\text{immediate}))$

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers; record the Boolean result of the comparison in GPR *rt*. If GPR *rs* is less than *immediate*, the result is 1 (true); otherwise, it is 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None

Operation:


```

if GPR[rs] < sign_extend(immediate) then
    GPR[rt] ← 0GPRLEN-1 || 1
else
    GPR[rt] ← 0GPRLEN
endif

```

Exceptions:

None



Format: SLTIU *rt*, *rs*, *immediate*

microMIPS

Purpose: Set on Less Than Immediate Unsigned

To record the result of an unsigned less-than comparison with a constant.

Description: $GPR[rt] \leftarrow (GPR[rs] < \text{sign_extend}(\text{immediate}))$

Compare the contents of GPR *rs* and the sign-extended 16-bit *immediate* as unsigned integers; record the Boolean result of the comparison in GPR *rt*. If GPR *rs* is less than *immediate*, the result is 1 (true); otherwise, it is 0 (false).

Because the 16-bit *immediate* is sign-extended before comparison, the instruction can represent the smallest or largest unsigned numbers. The representable values are at the minimum [0, 32767] or maximum [max_unsigned-32767, max_unsigned] end of the unsigned range.

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None

Operation:


```

if (0 || GPR[rs]) < (0 || sign_extend(immediate)) then
    GPR[rt] ← 0GPRLEN-1 || 1
else
    GPR[rt] ← 0GPRLEN
endif

```

Exceptions:

None



Format: SLTU rd, rs, rt

microMIPS

Purpose: Set on Less Than Unsigned

To record the result of an unsigned less-than comparison.

Description: $GPR[rd] \leftarrow (GPR[rs] < GPR[rt])$

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers; record the Boolean result of the comparison in GPR *rd*. If GPR *rs* is less than GPR *rt*, the result is 1 (true); otherwise, it is 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None

Operation:


```

if (0 || GPR[rs]) < (0 || GPR[rt]) then
    GPR[rd] ← 0GPRLEN-1 || 1
else
    GPR[rd] ← 0GPRLEN
endif

```

Exceptions:

None



Format: SQRT.fmt
 SQRT.S ft, fs
 SQRT.D ft, fs

MIPS32
MIPS32

Purpose: Floating Point Square Root

To compute the square root of an FP value.

Description: $FPR[ft] \leftarrow SQRT(FPR[fs])$

The square root of the value in FPR *fs* is calculated to infinite precision, rounded according to the current rounding mode in *FCSR*, and placed into FPR *ft*. The operand and result are values in format *fmt*.

If the value in FPR *fs* corresponds to -0 , the result is -0 .

Restrictions:

If the value in FPR *fs* is less than 0, an Invalid Operation condition is raised.

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(ft, fmt, SquareRoot(ValueFPR(fs, fmt)))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Inexact, Unimplemented Operation



Format: SRA *rt*, *rs*, *sa*

microMIPS

Purpose: Shift Word Right Arithmetic

To execute an arithmetic right-shift of a word by a fixed number of bits.

Description: $GPR[rt] \leftarrow GPR[rs] \gg sa$ (arithmetic)

The contents of the low-order 32-bit word of GPR *rs* are shifted right, duplicating the sign-bit (bit 31) in the emptied bits; the word result is sign-extended and placed in GPR *rt*. The bit-shift amount is specified by *sa*.

Restrictions:

On 64-bit processors, if GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
s ← sa
temp ← GPR[rs]31s || GPR[rs]31..s
GPR[rt] ← sign_extend(temp)

```

Exceptions:

None



Format: SRAV rd, rt, rs

microMIPS

Purpose: Shift Word Right Arithmetic Variable

To execute an arithmetic right-shift of a word by a variable number of bits.

Description: $GPR[rd] \leftarrow GPR[rt] \gg GPR[rs]$ (arithmetic)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, duplicating the sign-bit (bit 31) in the emptied bits; the word result is sign-extended and placed in GPR *rd*. The bit-shift amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
s ← GPR[rs]4..0
temp ← (GPR[rt]31)s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)

```

Exceptions:

None



Format: SRL *rt*, *rs*, *sa*

microMIPS

Purpose: Shift Word Right Logical

To execute a logical right-shift of a word by a fixed number of bits.

Description: $GPR[rt] \leftarrow GPR[rs] \gg sa$ (logical)

The contents of the low-order 32-bit word of GPR *rs* are shifted right, inserting zeros into the emptied bits. The word result is sign-extended and placed in GPR *rt*. The bit-shift amount is specified by *sa*.

Restrictions:

On 64-bit processors, if GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:


```

if NotWordValue(GPR[rs]) then
    UNPREDICTABLE
endif
s ← sa
temp ← 0s || GPR[rs]31..s
GPR[rt] ← sign_extend(temp)

```

Exceptions:

None



Format: SRLV rd, rt, rs

microMIPS

Purpose: Shift Word Right Logical Variable

To execute a logical right-shift of a word by a variable number of bits.

Description: $\text{GPR}[\text{rd}] \leftarrow \text{GPR}[\text{rt}] \gg \text{GPR}[\text{rs}]$ (logical)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, inserting zeros into the emptied bits; the word result is sign-extended and placed in GPR *rd*. The bit-shift amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
s ← GPR[rs]4..0
temp ← 0s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)

```

Exceptions:

None

31	26	25		11	10		6	5		0
POOL32A 000000	0 00000	0 00000	1 00001	0 0000			SLL32 000000			
6	5	5	5	5			6			

Format: SSNOP

microMIPS

Purpose: Superscalar No Operation

Break superscalar issue on a superscalar processor.

Description:

SSNOP is the assembly idiom used to denote superscalar no operation. The actual instruction is interpreted by the hardware as SLL r0, r0, 1.

This instruction alters the instruction issue behavior on a superscalar processor by forcing the SSNOP instruction to single-issue. The processor must then end the current instruction issue between the instruction previous to the SSNOP and the SSNOP. The SSNOP then issues alone in the next issue slot.

On a single-issue processor, this instruction is a NOP that takes an issue slot.

Restrictions:

None

Availability and Compatibility

Release 6: the special no-operation instruction SSNOP is deprecated: it behaves the same as a conventional NOP. Its special behavior with respect to instruction issue is no longer guaranteed. The EHB and JR.HB instructions are provided to clear execution and instruction hazards.

Assemblers targeting specifically Release 6 should reject the SSNOP instruction with an error.

Operation:

None

Exceptions:


None

Programming Notes:

SSNOP is intended for use primarily to allow the programmer control over CP0 hazards by converting instructions into cycles in a superscalar processor. For example, to insert at least two cycles between an MTC0 and an ERET, one would use the following sequence:

```
mtc0    x,y
ssnop
ssnop
eret
```

The MTC0 issues in cycle T. Because the SSNOP instructions must issue alone, they may issue no earlier than cycle T+1 and cycle T+2, respectively. Finally, the ERET issues no earlier than cycle T+3. Although the instruction after an SSNOP may issue no earlier than the cycle after the SSNOP is issued, that instruction may issue later. This is because other implementation-dependent issue rules may apply that prevent an issue in the next cycle. Processors should not introduce any unnecessary delay in issuing SSNOP instructions.



Format: SUB rd, rs, rt

microMIPS

Purpose: Subtract Word

To subtract 32-bit integers. If overflow occurs, then trap.

Description: $GPR[rd] \leftarrow GPR[rs] - GPR[rt]$

The 32-bit word value in GPR *rt* is subtracted from the 32-bit value in GPR *rs* to produce a 32-bit result. If the subtraction results in 32-bit 2's complement arithmetic overflow, then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 32-bit result is sign-extended and placed into GPR *rd*.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← (GPR[rs]31 | GPR[rs]31..0) - (GPR[rt]31 | GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← sign_extend(temp31..0)
endif


```

Exceptions:

Integer Overflow

Programming Notes:

SUBU performs the same arithmetic operation but does not trap on overflow.



Format: SUB.fmt
SUB.S fd, fs, ft
SUB.D fd, fs, ft

microMIPS
microMIPS

Purpose: Floating Point Subtract
To subtract FP values.

Description: $FPR[fd] \leftarrow FPR[fs] - FPR[ft]$
The value in FPR *ft* is subtracted from the value in FPR *fs*. The result is calculated to infinite precision, rounded according to the current rounding mode in *FCSR*, and placed into FPR *fd*. The operands and result are values in format *fmt*.


Restrictions:
The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If the fields are not valid, the result is **UNPREDICTABLE**.
The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

Availability and Compatibility:
SUB.PS has been removed in Release 6.

Operation:
`StoreFPR (fd, fmt, ValueFPR(fs, fmt) -fmt ValueFPR(ft, fmt))`

CPU Exceptions:
Coproprocessor Unusable, Reserved Instruction

FPU Exceptions:
Inexact, Overflow, Underflow, Invalid Op, Unimplemented Op



Format: SUBU *rd*, *rs*, *rt*

microMIPS

Purpose: Subtract Unsigned Word

To subtract 32-bit integers.

Description: $GPR[rd] \leftarrow GPR[rs] - GPR[rt]$

The 32-bit word value in GPR *rt* is subtracted from the 32-bit value in GPR *rs* and the 32-bit arithmetic result is sign-extended and placed into GPR *rd*.

No integer overflow exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* does not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```

if NotWordValue(GPR[rs]) or NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
temp ← GPR[rs] - GPR[rt]
GPR[rd] ← sign_extend(temp)


```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. It is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: SW *rt*, *offset*(*base*)

microMIPS

Purpose: Store Word

To store a word to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 32-bit word of GPR *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWE rt, offset(base)

microMIPS

Purpose: Store Word EVA

To store a word to user mode virtual address space when executing in kernel mode.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$

The least-significant 32-bit word of GPR *rt* is stored in memory at the location specified by the aligned effective address. The 9-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The SWE instruction functions the same as the SW instruction, except that address translation is performed using the user mode virtual address space mapping in the TLB when accessing an address within a memory segment configured to use the MUSUK access mode. Memory segments using UUSK or MUSK access modes are also accessible. Refer to Volume III, Enhanced Virtual Addressing section for additional information.

Implementation of this instruction is specified by the *Config5_{EVA}* field being set to 1.

Restrictions:

Only usable in kernel mode when accessing an address within a segment configured using UUSK, MUSK or MUSUK access mode.

Pre-Release 6: The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← GPR[rt]63-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch, Reserved Instruction, Coprocessor Unusable



SWC1 ft, offset(base)

microMIPS

Purpose: Store Word from Floating Point

To store a word from an FPR to memory.

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{FPR}[\text{ft}]$

The low 32-bit word from FPR *ft* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← ValueFPR(ft, UNINTERPRETED_WORD) || 08*bytesel
StoreMemory(CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWC2 rt, offset(base)

microMIPS

Purpose: Store Word from Coprocessor 2

To store a word from a COP2 register to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{CPR}[2, \text{rt}, 0]$

The low 32-bit word from COP2 (Coprocessor 2) register *rt* is stored in memory at the location specified by the aligned effective address. The signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

Pre-Release 6: An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Release 6 allows hardware to provide address misalignment support in lieu of requiring natural alignment.

Note: The pseudocode is not completely adapted for Release 6 misalignment support as the handling is implementation dependent.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
datadoubleword ← CPR[2,rt,0]63-8*bytesel..0 || 08*bytesel
StoreMemory(CCA, WORD, datadoubleword, pAddr, vAddr, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:

As shown in the instruction drawing above, Release 6 implements an 11-bit offset, whereas all release levels lower than Release 6 of the MIPS architecture implement a 16-bit offset.

31	26	25	16	15	6	5	0
POOL32A 000000	0 0000000000	type	SYNC 0110101101		POOL32AXf 111100		
6	5	5	10		6		

Format: SYNC (stype = 0 implied)
 SYNC stype

microMIPS
 microMIPS

Purpose: Synchronize Shared Memory

To order loads and stores for shared memory.

Release 6 (with *Config5_{GI}*=10/11) extends SYNC for Global Invalidate instructions (GINVI/GINVT).

Description:

These types of ordering guarantees are available through the SYNC instruction:

- Completion Barriers
- Ordering Barriers

Completion Barrier — Simple Description:

- The barrier affects only *uncached* and *cached coherent* loads and stores.
- The specified memory instructions (loads or stores or both) that occur before the SYNC instruction must be completed before the specified memory instructions after the SYNC are allowed to start.
- Loads are completed when the destination register is written. Stores are completed when the stored value is visible to every other processor in the system.

Completion Barrier — Detailed Description:

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must be already globally performed before any synchronizable specified memory instructions that occur after the SYNC are allowed to be performed, with respect to any other processor or coherent I/O module.
- The barrier does not guarantee the order in which instruction fetches are performed.
- A stype value of zero will always be defined such that it performs the most complete set of synchronization operations that are defined. This means stype zero always does a completion barrier that affects both loads and stores preceding the SYNC instruction and both loads and stores that are subsequent to the SYNC instruction. Non-zero values of stype may be defined by the architecture or specific implementations to perform synchronization behaviors that are less complete than that of stype zero. If an implementation does not use one of these non-zero values to define a different synchronization behavior, then that non-zero value of stype must act the same as stype zero completion barrier. This allows software written for an implementation with a lighter-weight barrier to work on another implementation which only implements the stype zero completion barrier.
- A completion barrier is required, potentially in conjunction with SSNOP (in Release 1 of the Architecture) or EHB (in Release 2 of the Architecture), to guarantee that memory reference results are visible across operating mode changes. For example, a completion barrier is required on some implementations on entry to and exit from Debug Mode to guarantee that memory effects are handled correctly.

SYNC behavior when the stype field is zero:

- A completion barrier that affects preceding loads and stores and subsequent loads and stores.

Ordering Barrier — Simple Description:

- The barrier affects only *uncached* and *cached coherent* loads and stores.
- The specified memory instructions (loads or stores or both) that occur before the SYNC instruction must always be ordered before the specified memory instructions after the SYNC.
- Memory instructions which are ordered before other memory instructions are processed by the load/store datapath first before the other memory instructions.

Ordering Barrier — Detailed Description:

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must reach a stage in the load/store datapath after which no instruction re-ordering is possible before any synchronizable specified memory instruction which occurs after the SYNC instruction in the instruction stream reaches the same stage in the load/store datapath.
- If any memory instruction before the SYNC instruction in program order, generates a memory request to the external memory and any memory instruction after the SYNC instruction in program order also generates a memory request to external memory, the memory request belonging to the older instruction must be globally performed before the time the memory request belonging to the younger instruction is globally performed.
- The barrier does not guarantee the order in which instruction fetches are performed.

As compared to the completion barrier, the ordering barrier is a lighter-weight operation as it does not require the specified instructions before the SYNC to be already completed. Instead it only requires that those specified instructions which are subsequent to the SYNC in the instruction stream are never re-ordered for processing ahead of the specified instructions which are before the SYNC in the instruction stream. This potentially reduces how many cycles the barrier instruction must stall before it completes.

The Acquire and Release barrier types are used to minimize the memory orderings that must be maintained and still have software synchronization work.

Implementations that do not use any of the non-zero values of stype to define different barriers, such as ordering barriers, must make those stype values act the same as stype zero.

For the purposes of this description, the CACHE, PREF and PREFX instructions are treated as loads and stores. That is, these instructions and the memory transactions sourced by these instructions obey the ordering and completion rules of the SYNC instruction.

If Global Invalidate instructions are supported in Release 6, then SYNC (stype=0x14) acts as a completion barrier to ensure completion of any preceding GINVI or GINVT operation. This SYNC operation is globalized and only completes if all preceding GINVI or GINVT operations related to the same program have completed in the system. (Any references to GINVT also imply GINVTG, available in a virtualized MIPS system. GINVT however will be used exclusively.)

A system that implements the Global Invalidates also requires that the completion of this SYNC be constrained by legacy SYNCI operations. Thus SYNC (stype=0x14) can also be used to determine whether preceding (in program order) SYNCI operations have completed.

The SYNC (stype=0x14) also act as an ordering barrier as described in [Table 10.28](#).

In the typical use cases, a single GINVI is used by itself to invalidate caches and would be followed by a SYNC (stype=0x14).

In the case of GINVT, multiple GINVT could be used to invalidate multiple TLB mappings, and the SYNC (stype=0x14) would be used to guaranteed completion of any number of GINVTs preceding it.

[Table 10.28](#) lists the available completion barrier and ordering barriers behaviors that can be specified using the stype field.

Table 10.28 Encodings of the Bits[10:6] of the SYNC instruction; the SType Field

Code	Name	Older instructions which must reach the load/store ordering point before the SYNC instruction completes.	Younger instructions which must reach the load/store ordering point only after the SYNC instruction completes.	Older instructions which must be globally performed when the SYNC instruction completes	Compliance
0x0	SYNC or SYNC 0	Loads, Stores	Loads, Stores	Loads, Stores	Required
0x4	SYNC_WMB or SYNC 4	Stores	Stores		Optional
0x10	SYNC_MB or SYNC 16	Loads, Stores	Loads, Stores		Optional
0x11	SYNC_ACQUIRE or SYNC 17	Loads	Loads, Stores		Optional
0x12	SYNC_RELEASE or SYNC 18	Loads, Stores	Stores		Optional
0x13	SYNC_RMB or SYNC 19	Loads	Loads		Optional
0x1-0x3, 0x5-0xF					Implementation-Specific and Vendor Specific Sync Types
0x14	SYNC_GINV	Loads, Stores	Loads, Stores	GINVI, GINVT, SYNCI	Release 6 w/ <i>Config5_{GI}</i> = 10/11 otherwise Reserved
0x15 - 0x1F	RESERVED				Reserved for MIPS Technologies for future extension of the architecture.

Terms:

Synchronizable: A load or store instruction is *synchronizable* if the load or store occurs to a physical location in shared memory using a virtual location with a memory access type of either *uncached* or *cached coherent*. *Shared memory* is memory that can be accessed by more than one processor or by a coherent I/O system module.

Performed load: A load instruction is *performed* when the value returned by the load has been determined. The result of a load on processor A has been *determined* with respect to processor or coherent I/O module B when a subsequent store to the location by B cannot affect the value returned by the load. The store by B must use the same memory access type as the load.

Performed store: A store instruction is *performed* when the store is observable. A store on processor A is *observable* with respect to processor or coherent I/O module B when a subsequent load of the location by B returns the value written by the store. The load by B must use the same memory access type as the store.

Globally performed load: A load instruction is *globally performed* when it is performed with respect to all processors and coherent I/O modules capable of storing to the location.

Globally performed store: A store instruction is *globally performed* when it is globally observable. It is *globally observable* when it is observable by all processors and I/O modules capable of loading from the location.

Coherent I/O module: A *coherent I/O module* is an Input/Output system component that performs coherent Direct Memory Access (DMA). It reads and writes memory independently as though it were a processor doing loads and stores to locations with a memory access type of *cached coherent*.

Load/Store Datapath: The portion of the processor which handles the load/store data requests coming from the processor pipeline and processes those requests within the cache and memory system hierarchy.

Restrictions:

The effect of SYNC on the global order of loads and stores for memory access types other than *uncached* and *cached coherent* is **UNPREDICTABLE**.

Operation:

```
SyncOperation(stype)
```

Exceptions:

None

Programming Notes:

A processor executing load and store instructions observes the order in which loads and stores using the same memory access type occur in the instruction stream; this is known as *program order*.

A *parallel program* has multiple instruction streams that can execute simultaneously on different processors. In multiprocessor (MP) systems, the order in which the effects of loads and stores are observed by other processors—the *global order* of the loads and store—determines the actions necessary to reliably share data in parallel programs.

When all processors observe the effects of loads and stores in program order, the system is *strongly ordered*. On such systems, parallel programs can reliably share data without explicit actions in the programs. For such a system, SYNC has the same effect as a NOP. Executing SYNC on such a system is not necessary, but neither is it an error.

If a multiprocessor system is not strongly ordered, the effects of load and store instructions executed by one processor may be observed out of program order by other processors. On such systems, parallel programs must take explicit actions to reliably share data. At critical points in the program, the effects of loads and stores from an instruction stream must occur in the same order for all processors. SYNC separates the loads and stores executed on the processor into two groups, and the effect of all loads and stores in one group is seen by all processors before the effect of any load or store in the subsequent group. In effect, SYNC causes the system to be strongly ordered for the executing processor at the instant that the SYNC is executed.

Many MIPS-based multiprocessor systems are strongly ordered or have a mode in which they operate as strongly ordered for at least one memory access type. The MIPS architecture also permits implementation of MP systems that are not strongly ordered; SYNC enables the reliable use of shared memory on such systems. A parallel program that does not use SYNC generally does not operate on a system that is not strongly ordered. However, a program that does use SYNC works on both types of systems. (System-specific documentation describes the actions needed to reliably share data in parallel programs for that system.)

The behavior of a load or store using one memory access type is **UNPREDICTABLE** if a load or store was previously made to the same physical location using a different memory access type. The presence of a SYNC between the references does not alter this behavior.

SYNC affects the order in which the effects of load and store instructions appear to all processors; it does not generally affect the physical memory-system ordering or synchronization issues that arise in system programming. The effect of SYNC on implementation-specific aspects of the cached memory system, such as writeback buffers, is not defined.

```
# Processor A (writer)
# Conditions at entry:
# The value 0 has been stored in FLAG and that value is observable by B
SW    R1, DATA      # change shared DATA value
LI    R2, 1
SYNC                      # Perform DATA store before performing FLAG store
SW    R2, FLAG        # say that the shared DATA value is valid

# Processor B (reader)
LI    R2, 1
1: LW  R1, FLAG # Get FLAG
    BNE R2, R1, 1B# if it says that DATA is not valid, poll again
    NOP
    SYNC                # FLAG value checked before doing DATA read
    LW  R1, DATA # Read (valid) shared DATA value
```

The code fragments above shows how SYNC can be used to coordinate the use of shared data between separate writer and reader instruction streams in a multiprocessor environment. The FLAG location is used by the instruction streams to determine whether the shared data item DATA is valid. The SYNC executed by processor A forces the store of DATA to be performed globally before the store to FLAG is performed. The SYNC executed by processor B ensures that DATA is not read until after the FLAG value indicates that the shared data is valid.

POOL32I 010000	SYNCI 01100	base	offset
6	5	5	16

Format: SYNCI offset (base)

microMIPS

Purpose: Synchronize Caches to Make Instruction Writes Effective

To synchronize all caches to make instruction writes effective.

Description:

This instruction is used after a new instruction stream is written to make the new instructions effective relative to an instruction fetch, when used in conjunction with the SYNC and JALR.HB, JR.HB, or ERET instructions, as described below. Unlike the CACHE instruction, the SYNCI instruction is available in all operating modes in an implementation of Release 2 of the architecture.

The 16-bit offset is sign-extended and added to the contents of the base register to form an effective address. The effective address is used to address the cache line in all caches which may need to be synchronized with the write of the new instructions. The operation occurs only on the cache line which may contain the effective address. One SYNCI instruction is required for every cache line that was written. See the Programming Notes below.

A TLB Refill and TLB Invalid (both with cause code equal TLBL) exception can occur as a by product of this instruction. This instruction never causes TLB Modified exceptions nor TLB Refill exceptions with a cause code of TLBS. This instruction never causes Execute-Inhibit nor Read-Inhibit exceptions.

A Cache Error exception may occur as a by product of this instruction. For example, if a writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Similarly, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error.

An Address Error Exception (with cause code equal AdEL) may occur if the effective address references a portion of the kernel address space which would normally result in such an exception. It is implementation dependent whether such an exception does occur.

It is implementation dependent whether a data watch is triggered by a SYNCI instruction whose address matches the Watch register address match conditions.

Restrictions:

The operation of the processor is **UNPREDICTABLE** if the effective address references any instruction cache line that contains instructions to be executed between the SYNCI and the subsequent JALR.HB, JR.HB, or ERET instruction required to clear the instruction hazard.

The SYNCI instruction has no effect on cache lines that were previously locked with the CACHE instruction. If correct software operation depends on the state of a locked line, the CACHE instruction must be used to synchronize the caches.

Full visibility of the new instruction stream requires execution of a subsequent SYNC instruction, followed by a JALR.HB, JR.HB, DERET, or ERET instruction. The operation of the processor is **UNPREDICTABLE** if this sequence is not followed.

SYNCI globalization:

The SYNCI instruction acts on the current processor at a minimum. Implementations are required to affect caches outside the current processor to perform the operation on the current processor (as might be the case if multiple processors share an L2 or L3 cache).

In multiprocessor implementations where instruction caches are coherently maintained by hardware, the SYNCI instruction should behave as a NOP instruction.

In multiprocessor implementations where instruction caches are not coherently maintained by hardware, the SYNCI instruction may optionally affect all coherent icaches within the system. If the effective address uses a coherent Cacheability and Coherency Attribute (CCA), then the operation may be *globalized*, meaning it is broadcast to all of the coherent instruction caches within the system. If the effective address does not use one of the coherent CCAs, there is no broadcast of the SYNCI operation. If multiple levels of caches are to be affected by one SYNCI instruction, all of the affected cache levels must be processed in the same manner - either all affected cache levels use the globalized behavior or all affected cache levels use the non-globalized behavior.

Pre-Release 6: Portable software could not rely on the optional *globalization* of SYNCI. Strictly portable software without implementation specific awareness could only rely on expensive “instruction cache shutdown” using inter-processor interrupts.

Release 6: SYNCI *globalization* is required. Compliant implementations must globalize SYNCI, and portable software can rely on this behavior.

Availability and Compatibility

This instruction has been recoded for Release 6.

Operation:

```
vaddr ← GPR[base] + sign_extend(offset)
SynchronizeCacheLines(vaddr)      /* Operate on all caches */
```

Exceptions:

Reserved Instruction exception (Release 1 implementations only)

TLB Refill Exception

TLB Invalid Exception

Address Error Exception

Cache Error Exception

Bus Error Exception


Programming Notes:

When the instruction stream is written, the SYNCI instruction should be used in conjunction with other instructions to make the newly-written instructions effective. The following example shows a routine which can be called after the new instruction stream is written to make those changes effective. The SYNCI instruction could be replaced with the corresponding sequence of CACHE instructions (when access to Coprocessor 0 is available), and that the JR.HB instruction could be replaced with JALR.HB, ERET, or DERET instructions, as appropriate. A SYNC instruction is required between the final SYNCI instruction in the loop and the instruction that clears instruction hazards.

```
/*
 * This routine makes changes to the instruction stream effective to the
 * hardware. It should be called after the instruction stream is written.
 * On return, the new instructions are effective.
 *
 * Inputs:
 *   a0 = Start address of new instruction stream
 *   a1 = Size, in bytes, of new instruction stream
 */

beq    a1, zero, 20f      /* If size==0, */
nop                    /* branch around */
addu   a1, a0, a1        /* Calculate end address + 1 */
                        /* (daddu for 64-bit addressing) */
rdhwr  v0, HW_SYNCI_Step /* Get step size for SYNCI from new */
                        /* Release 2 instruction */
```

```
    beq    v0, zero, 20f    /* If no caches require synchronization, */
    nop                                /*   branch around */
10: synci  0(a0)            /* Synchronize all caches around address */
    addu   a0, a0, v0        /* Add step size in delay slot */
                                /* (daddu for 64-bit addressing) */
    sltu   v1, a0, a1        /* Compare current with end address */
    bne    v1, zero, 10b     /* Branch if more to do */
    nop                                /*   branch around */
    sync                                /* Clear memory hazards */
20: jr.hb  ra                /* Return, clearing instruction hazards */
    nop
```

Format: SYSCALL

microMIPS

Purpose: System Call

To cause a System Call exception.

Description:

A system call exception occurs, immediately and unconditionally transferring control to the exception handler.

The *code* field is available for use as software parameters, but may be retrieved by the exception handler by loading the contents of the memory word containing the instruction. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

`SignalException(SystemCall)`

Exceptions:

System Call



Format: TEQ *rs*, *rt*

microMIPS

Purpose: Trap if Equal

To compare GPRs and do a conditional trap.

Description: if GPR[*rs*] = GPR[*rt*] then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers. If GPR *rs* is equal to GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software may load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

```
if GPR[rs] = GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TGE *rs*, *rt*

microMIPS

Purpose: Trap if Greater or Equal

To compare GPRs and do a conditional trap.

Description: if $GPR[rs] \geq GPR[rt]$ then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers. If GPR *rs* is greater than or equal to GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, the system software may load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

```
if GPR[rs] ≥ GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TGEU *rs*, *rt*

microMIPS

Purpose: Trap if Greater or Equal Unsigned

To compare GPRs and do a conditional trap.

Description: if $GPR[rs] \geq GPR[rt]$ then Trap

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers. If GPR *rs* is greater than or equal to GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, the system software may load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:

None

Operation:

```
if (0 || GPR[rs]) ≥ (0 || GPR[rt]) then
    SignalException(Trap)
endif
```

Exceptions:

Trap

31	26	25	16	15	6	5	0
POOL32A 000000		0000000000		TLBINV 0100001101		POOL32Axf 111100	
6		10		10		6	

Format: TLBINV

microMIPS

Purpose: TLB Invalidate

TLBINV invalidates a set of TLB entries based on ASID and Index match. The virtual address is ignored in the entry match. TLB entries which have their G bit set to 1 are not modified.

Implementation of the TLBINV instruction is optional. The implementation of this instruction is indicated by the IE field in *Config4*.

Support for TLBINV is recommend for implementations supporting VTLB/FTLB type of MMU.

Implementation of *EntryHI_{EHINV}* field is required for implementation of TLBINV instruction.

Description:

On execution of the TLBINV instruction, the set of TLB entries with matching ASID are marked invalid, excluding those TLB entries which have their G bit set to 1.

The *EntryHI_{ASID}* field has to be set to the appropriate ASID value before executing the TLBINV instruction.

Behavior of the TLBINV instruction applies to all applicable TLB entries and is unaffected by the setting of the *Wired* register.

- For JTLB-based MMU (*Config_{MT}*=1):

All matching entries in the JTLB are invalidated. The *Index* register is unused.

- For VTLB/FTLB -based MMU (*Config_{MT}*=4):

If TLB invalidate walk is implemented in software (*Config_{4IE}*=2), then software must do these steps to flush the entire MMU:

- one TLBINV instruction is executed with an index in VTLB range (invalidates all matching VTLB entries)
- a TLBINV instruction is executed for each FTLB set (invalidates all matching entries in FTLB set)

If TLB invalidate walk is implemented in hardware (*Config_{4IE}*=3), then software must do these steps to flush the entire MMU:

- one TLBINV instruction is executed (invalidates all matching entries in both FTLB & VTLB). In this case, *Index* is unused.

Restrictions:

When *Config_{4MT}* = 4 and *Config_{4IE}* = 2, the operation is **UNDEFINED** if the contents of the *Index* register are greater than or equal to the number of available TLB entries.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Availability and Compatibility:

Implementation of the TLBINV instruction is optional. The implementation of this instruction is indicated by the IE

field in *Config4*.

Implementation of *EntryHi_{EHINV}* field is required for implementation of TLBINV instruction.

Pre-Release 6, support for TLBINV is recommended for implementations supporting VTLB/FTLB type of MMU. Release 6 (and subsequent releases) support for TLBINV is required for implementations supporting VTLB/FTLB type of MMU.

Release 6: On processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU (*Config_{MT}* = 2 or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```

if ( ConfigMT=1 or (ConfigMT=4 & Config4IE=2 & Index < VTLBsize() ) )
    startnum ← 0
    endnum ← VTLBsize() - 1
endif
// treating VTLB and FTLB as one array
if (ConfigMT=4 & Config4IE=2 & Index ≥ VTLBsize(); )
    startnum ← start of selected FTLB set // implementation specific
    endnum ← end of selected FTLB set - 1 //implementation specific
endif

if (ConfigMT=4 & Config4IE=3))
    startnum ← 0
    endnum ← VTLBsize() + FTLBsize() - 1;
endif

for (i = startnum to endnum)
    if (TLB[i]ASID = EntryHiASID & TLB[i]G = 0)
        TLB[i]VPN2_invalid ← 1
    endif
endfor

function VTLBsize
    SizeExt = ArchRev() ≥ 6                ? Config4VTLBSizeExt
        : Config4MMUExtDef == 3            ? Config4VTLBSizeExt
        : Config4MMUExtDef == 1            ? Config4MMUSizeExt
        :                                0
    ;
    return 1 + ( (SizeExt << 6) | Config1MMUSize );
endfunction

function FTLBsize
    if ( Config1MT == 4 ) then
        return ( Config4FTLBWays + 2 ) * ( 1 << C0_Config4FTLBsets );
    else
        return 0;
    endif
endfunction

```

Exceptions:

Coprocessor Unusable,

31	26	25	16	15	6	5	0
POOL32A 000000		0000000000		TLBINV 0101001101		POOL32Axf 111100	
6		10		10		6	

Format: TLBINVF

microMIPS

Purpose: TLB Invalidate Flush

TLBINVF invalidates a set of TLB entries based on *Index* match. The virtual address and ASID are ignored in the entry match.

Implementation of the TLBINVF instruction is optional. The implementation of this instruction is indicated by the IE field in *Config4*.

Support for TLBINVF is recommend for implementations supporting VTLB/FTLB type of MMU.

Implementation of the *EntryHI_{EHINV}* field is required for implementation of TLBINV and TLBINVF instructions.

Description:

On execution of the TLBINVF instruction, all entries within range of *Index* are invalidated.

Behavior of the TLBINVF instruction applies to all applicable TLB entries and is unaffected by the setting of the *Wired* register.

- For JTLB-based MMU (*Config_{MT}*=1):

TLBINVF causes all entries in the JTLB to be invalidated. *Index* is unused.

- For VTLB/FTLB-based MMU (*Config_{MT}*=4):

If TLB invalidate walk is implemented in your software (*Config4_{IE}*=2), then your software must do these steps to flush the entire MMU:

- one TLBINVF instruction is executed with an index in VTLB range (invalidates all VTLB entries)
- a TLBINVF instruction is executed for each FTLB set (invalidates all entries in FTLB set)

If TLB invalidate walk is implemented in hardware (*Config4_{IE}*=3), then software must do these steps to flush the entire MMU:

- one TLBINVF instruction is executed (invalidates all entries in both FTLB & VTLB). In this case, *Index* is unused.

Restrictions:

When *Config_{MT}*=4 and *Config_{IE}*=2, the operation is **UNDEFINED** if the contents of the *Index* register are greater than or equal to the number of available TLB entries.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Availability and Compatibility:

Implementation of the TLBINVF instruction is optional. The implementation of this instruction is indicated by the IE field in *Config4*.

Implementation of *EntryHI_{EHINV}* field is required for implementation of TLBINVF instruction.

Pre-Release 6, support for TLBINVF is recommended for implementations supporting VTLB/FTLB type of MMU. Release 6 (and subsequent releases) support for TLBINV is required for implementations supporting VTLB/FTLB type of MMU.

Release 6: On processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU (*Config_{MT}* = 2 or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```

if ( ConfigMT=1 or (ConfigMT=4 & Config4IE=2 & Index < VTLBsize() ))
    startnum ← 0
    endnum ← VTLBsize() - 1
endif
// treating VTLB and FTLB as one array
if (ConfigMT=4 & Config4IE=2 & Index ≥ VTLBsize(); )
    startnum ← start of selected FTLB set // implementation specific
    endnum ← end of selected FTLB set - 1 //implementation specific
endif

if (ConfigMT=4 & Config4IE=3))
    startnum ← 0
    endnum ← TLBsize() + FTLBsize() - 1;
endif

for (i = startnum to endnum)
    TLB[i]VPN2_invalid ← 1
endfor

function VTLBsize
    SizeExt = ArchRev() ≥ 6           ? Config4VTLBSizeExt
    : Config4MMUExtDef == 3           ? Config4VTLBSizeExt
    : Config4MMUExtDef == 1           ? Config4MMUSizeExt
    :                                0
    ;
    return 1 + ( (SizeExt << 6) | Config1MMUSize );
endfunction

function FTLBsize
    if ( Config1MT == 4 ) then
        return ( Config4FTLBWays + 2 ) * ( 1 << C0_Config4FTLBsets );
    else
        return 0;
    endif
endfunction

```

Exceptions:

Coprocessor Unusable,

31	26	25	16	15	6	5	0
POOL32A 000000	0 0000000000					TLBP 0000001101	POOL32AXf 111100
6	10					10	6

Format: TLBP

microMIPS

Purpose: Probe TLB for Matching Entry

To find a matching entry in the TLB.

Description:

The *Index* register is loaded with the address of the TLB entry whose contents match the contents of the *EntryHi* register. If no TLB entry matches, the high-order bit of the *Index* register is set.

- In Release 1 of the Architecture, it is implementation dependent whether multiple TLB matches are detected on a TLBP. However, implementations are strongly encouraged to report multiple TLB matches only on a TLB write.
- In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write.
- In Release 3 of the Architecture, multiple TLB matches may be reported on either TLB write or TLB probe.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Release 6: Processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU (*Config_{MT}* = 2 or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```

Index ← 1 || UNPREDICTABLE31
for i in 00 ... TLBEntries-1
    if ((TLB[i]VPN2 and not (TLB[i]Mask)) =
        (EntryHiVPN2 and not (TLB[i]Mask))) and
        (TLB[i]R = EntryHiR) and
        ((TLB[i]G = 1) or (TLB[i]ASID = EntryHiASID)) then
        Index ← i
    endif
endfor

```

Exceptions:

Coprocessor Unusable, Machine Check

31	26	25	16	15	6	5	0
POOL32A 000000	0 0000000000		TLBR 0001001101		POOL32AXf 111100		
6	10		10		6		

Format: TLBR

microMIPS

Purpose: Read Indexed TLB Entry

To read an entry from the TLB.

Description:

The *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers are loaded with the contents of the TLB entry pointed to by the *Index* register.

- In Release 1 of the Architecture, it is implementation dependent whether multiple TLB matches are detected on a TLBR. However, implementations are strongly encouraged to report multiple TLB matches only on a TLB write.
- In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write.
- In Release 3 of the Architecture, multiple TLB matches may be detected on a TLBR.

In an implementation supporting TLB entry invalidation (*Config4_{IE}* ≥ 1), reading an invalidated TLB entry causes *EntryLo0* and *EntryLo1* to be set to 0, *EntryHi_{EHINV}* to be set to 1, all other *EntryHi* bits to be set to 0, and *PageMask* to be set to a value representing the minimum supported page size..

The value written to the *EntryHi*, *EntryLo0*, and *EntryLo1* registers may be different from the original written value to the TLB via these registers in that:

- The value returned in the *VPN2* field of the *EntryHi* register may have those bits set to zero corresponding to the one bits in the Mask field of the TLB entry (the least-significant bit of *VPN2* corresponds to the least-significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed after a TLB entry is written and then read.
- The value returned in the *PFN* field of the *EntryLo0* and *EntryLo1* registers may have those bits set to zero corresponding to the one bits in the Mask field of the TLB entry (the least significant bit of *PFN* corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed after a TLB entry is written and then read.
- The value returned in the *G* bit in both the *EntryLo0* and *EntryLo1* registers comes from the single *G* bit in the TLB entry. Recall that this bit was set from the logical AND of the two *G* bits in *EntryLo0* and *EntryLo1* when the TLB was written.

Restrictions:

The operation is **UNDEFINED** if the contents of the Index register are greater than or equal to the number of TLB entries in the processor.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Release 6: Processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU (*Config_{MT}* = 2 or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```

i ← Index
if i > (TLBEntries - 1) then
    UNDEFINED
endif

```

```

if ( (Config4IB ≥ 1) and TLB[i]VPN2_invalid = 1) then
    PagemaskMask ← 0 // or value representing minimum page size
    EntryHi ← 0
    EntryLo1 ← 0
    EntryLo0 ← 0
    EntryHiEHINV ← 1
else
    PageMaskMask ← TLB[i]Mask
    EntryHi ← TLB[i]R || 0Fill ||
        (TLB[i]VPN2 and not TLB[i]Mask) || # Masking implem dependent
        05 || TLB[i]ASID
    EntryLo1 ← 0Fill ||
        (TLB[i]PFN1 and not TLB[i]Mask) || # Masking mplem dependent
        TLB[i]C1 || TLB[i]D1 || TLB[i]V1 || TLB[i]G
    EntryLo0 ← 0Fill ||
        (TLB[i]PFN0 and not TLB[i]Mask) || # Masking mplem dependent
        TLB[i]C0 || TLB[i]D0 || TLB[i]V0 || TLB[i]G
endif

```

Exceptions:

Coprocessor Unusable, Machine Check

31	26	25	16	15	6	5	0
POOL32A 000000	0000000000					TLBWI 0010001101	POOL32Axf 111100
6	10					10	6

Format: TLBWI

microMIPS

Purpose: Write Indexed TLB EntryTo write or invalidate a TLB entry indexed by the *Index* register.**Description:**If $Config4_{IE} == 0$ or $EntryHi_{EHINV} = 0$:

The TLB entry pointed to by the Index register is written from the contents of the *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers. It is implementation dependent whether multiple TLB matches are detected on a TLBWI. In such an instance, a Machine Check Exception is signaled.

In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write. The information written to the TLB entry may be different from that in the *EntryHi*, *EntryLo0*, and *EntryLo1* registers, in that:

- The value written to the VPN2 field of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of the *PageMask* register (the least significant bit of VPN2 corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The value written to the PFN0 and PFN1 fields of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of *PageMask* register (the least significant bit of PFN corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The single G bit in the TLB entry is set from the logical AND of the G bits in the *EntryLo0* and *EntryLo1* registers.

If $Config4_{IE} \geq 1$ and $EntryHi_{EHINV} = 1$:

The TLB entry pointed to by the Index register has its VPN2 field marked as invalid. This causes the entry to be ignored on TLB matches for memory accesses. No Machine Check is generated.

Restrictions:

The operation is **UNDEFINED** if the contents of the Index register are greater than or equal to the number of TLB entries in the processor.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Release 6: Processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU ($Config_{MT} = 2$ or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```

i ← Index
if (Config4IE ≥ 1) then
    TLB[i]VPN2_invalid ← 0
    if ( EntryHIEHINV=1 ) then

```

```

        TLB[i]VPN2_invalid ← 1
        break
    endif
endif
TLB[i]Mask ← PageMaskMask
TLB[i]R ← EntryHiR
TLB[i]VPN2 ← EntryHiVPN2 and not PageMaskMask # Implementation dependent
TLB[i]ASID ← EntryHiASID
TLB[i]G ← EntryLo1G and EntryLo0G
TLB[i]PFN1 ← EntryLo1PFN and not PageMaskMask # Implementation dependent
TLB[i]C1 ← EntryLo1C
TLB[i]D1 ← EntryLo1D
TLB[i]V1 ← EntryLo1V
TLB[i]PFN0 ← EntryLo0PFN and not PageMaskMask # Implementation dependent
TLB[i]C0 ← EntryLo0C
TLB[i]D0 ← EntryLo0D
TLB[i]V0 ← EntryLo0V

```

Exceptions:

Coprocessor Unusable, Machine Check

31	26	25	16	15	6	5	0
POOL32A 000000	0000000000				TLBWR 0011001101	POOL32Axf 111100	
6	10				10	6	

Format: TLBWR

microMIPS

Purpose: Write Random TLB Entry

To write a TLB entry indexed by the *Random* register, or, in Release 6, write a TLB entry indexed by an implementation-defined location.

Description:

The TLB entry pointed to by the *Random* register is written from the contents of the *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers. It is implementation dependent whether multiple TLB matches are detected on a TLBWR. In such an instance, a Machine Check Exception is signaled.

In Release 6, the *Random* register has been removed. References to *Random* refer to an implementation-determined value that is not visible to software.

In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write. The information written to the TLB entry may be different from that in the *EntryHi*, *EntryLo0*, and *EntryLo1* registers, in that:

- The value written to the VPN2 field of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of the *PageMask* register (the least significant bit of VPN2 corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The value written to the PFN0 and PFN1 fields of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of *PageMask* register (the least significant bit of PFN corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The single G bit in the TLB entry is set from the logical AND of the G bits in the *EntryLo0* and *EntryLo1* registers.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Release 6: Processors that include a Block Address Translation (BAT) or Fixed Mapping (FM) MMU (*Config_{MT}* = 2 or 3), the operation of this instruction causes a Reserved Instruction exception (RI).

Operation:

```


i ← Random
if (Config4IE ≥ 1) then
    TLB[i]VPN2_invalid ← 0
endif
TLB[i]Mask ← PageMaskMask
TLB[i]R ← EntryHiR
TLB[i]VPN2 ← EntryHiVPN2 and not PageMaskMask # Implementation dependent
TLB[i]ASID ← EntryHiASID
TLB[i]G ← EntryLo1G and EntryLo0G
TLB[i]PFN1 ← EntryLo1PFN and not PageMaskMask # Implementation dependent
TLB[i]C1 ← EntryLo1C
TLB[i]D1 ← EntryLo1D
TLB[i]V1 ← EntryLo1V

```


$TLB[i]_{PFN0} \leftarrow EntryLo0_{PFN} \text{ and not } PageMask_{Mask} \text{ \# Implementation dependent}$
 $TLB[i]_{C0} \leftarrow EntryLo0_C$
 $TLB[i]_{D0} \leftarrow EntryLo0_D$
 $TLB[i]_{V0} \leftarrow EntryLo0_V$

Exceptions:

Coprocessor Unusable, Machine Check



Format: TLT *rs*, *rt*

microMIPS

Purpose: Trap if Less Than

To compare GPRs and do a conditional trap.

Description: if GPR[*rs*] < GPR[*rt*] then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers. If GPR *rs* is less than GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

```
if GPR[rs] < GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TLTU *rs*, *rt*

microMIPS

Purpose: Trap if Less Than Unsigned

To compare GPRs and do a conditional trap.

Description: if GPR[*rs*] < GPR[*rt*] then Trap

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers. If GPR *rs* is less than GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

```
if (0 || GPR[rs]) < (0 || GPR[rt]) then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TNE *rs*, *rt*

microMIPS

Purpose: Trap if Not Equal

To compare GPRs and do a conditional trap.

Description: if GPR[*rs*] \neq GPR[*rt*] then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers. If GPR *rs* is not equal to GPR *rt*, then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory. Alternatively, if CP0 *BadInstr* is implemented, the *code* field may be obtained from *BadInstr*.

Restrictions:


None

Operation:

```
if GPR[rs]  $\neq$  GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TRUNC.L.fmt
 TRUNC.L.S ft, fs
 TRUNC.L.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Truncate to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding toward zero.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long-fixed point format and rounded toward zero (rounding mode 1). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. In this case the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *ft* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{63}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{63}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{63}-1$
- $-2^{63}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{63}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *fd* for long fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in the *FR*=0 32-bit FPU register model; it is predictable if executing on a 64-bit FPU in the *FR*=1 mode, but not with *FR*=0, and not on a 32-bit FPU.

Operation:


$\text{StoreFPR}(ft, L, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Inexact



Format: TRUNC.W.fmt
 TRUNC.W.S ft, fs
 TRUNC.W.D ft, fs

microMIPS
 microMIPS

Purpose: Floating Point Truncate to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding toward zero.

Description: $FPR[ft] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format using rounding toward zero (rounding mode 1). The result is placed in FPR *ft*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. In this case the Invalid Operation flag is set in the *FCSR*. If the Invalid Operation *Enable* bit is set in the *FCSR*, no result is written to *ft* and an Invalid Operation exception is taken immediately. Otherwise, a default result is written to *ft*. On cores with $FCSR_{NAN2008}=0$, the default result is $2^{31}-1$. On cores with $FCSR_{NAN2008}=1$, the default result is:

- 0 when the input value is NaN
- $2^{31}-1$ when the input value is $+\infty$ or rounds to a number larger than $2^{31}-1$
- $-2^{31}-1$ when the input value is $-\infty$ or rounds to a number smaller than $-2^{31}-1$

Restrictions:

The fields *fs* and *ft* must specify valid FPRs: *fs* for type *fmt* and *fd* for word fixed point. If the fields are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:

`StoreFPR(ft, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Invalid Operation, Unimplemented Operation

31	26	25	16	15	6	5	0
POOL32A 000000	Implementation-dependent code					WAIT 1001001101	POOL32AXf 111100
6	10					10	6

Format: WAIT

microMIPS

Purpose: Enter Standby Mode

Wait for Event

Description:

The WAIT instruction performs an implementation-dependent operation, involving a lower power mode. Software may use the code bits of the instruction to communicate additional information to the processor. The processor may use this information as control for the lower power mode. A value of zero for code bits is the default and must be valid in all implementations.

The WAIT instruction is implemented by stalling the pipeline at the completion of the instruction and entering a lower power mode. The pipeline is restarted when an external event, such as an interrupt or external request occurs, and execution continues with the instruction following the WAIT instruction. It is implementation-dependent whether the pipeline restarts when a non-enabled interrupt is requested. In this case, software must poll for the cause of the restart. The assertion of any reset or NMI must restart the pipeline and the corresponding exception must be taken.

If the pipeline restarts as the result of an enabled interrupt, that interrupt is taken between the WAIT instruction and the following instruction (EPC for the interrupt points at the instruction following the WAIT instruction).

In Release 6, the behavior of WAIT has been modified to make it a requirement that a processor that has disabled operation as a result of executing a WAIT will resume operation on arrival of an interrupt even if interrupts are not enabled.

In Release 6, the encoding of WAIT with bits 25:16 of the opcode set to 0 will never disable CP0 *Count* on an active WAIT instruction. In particular, this modification has been added to architecturally specify that CP0 *Count* is not disabled on execution of WAIT with default code of 0. Prior to Release 6, whether *Count* is disabled was implementation-dependent. In the future, other encodings of WAIT may be defined which specify other forms of power-saving or stand-by modes. If not implemented, then such unimplemented encodings must default to WAIT 0.

Restrictions:

Pre-Release 6: The operation of the processor is **UNDEFINED** if a WAIT instruction is executed in the delay slot of a branch or jump instruction.

Release 6: Implementations are required to signal a Reserved Instruction exception if WAIT is encountered in the delay slot or forbidden slot of a branch or jump instruction.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Operation:

Pre-Release 6:

```
I: Enter implementation dependent lower power mode
I+1:/* Potential interrupt taken here */
```

Release 6:

```
I: if IsCoprocessorEnabled(0) then
    while ( !interrupt_pending_and_not_masked_out() &&
            !implementation_dependent_wake_event() )
        < enter or remain in low power mode or stand-by mode>
```

```

    else
        SignalException(CoprocessorUnusable, 0)
    endif

I+1:  if ( interrupt_pending() && interrupts_enabled() ) then
        EPC ← PC + 4
        < process interrupt; execute ERET eventually >
    else
        // unblock on non-enabled interrupt or imp dep wake event.
        PC ← PC + 4
        < continue execution at instruction after wait >
    endif

function interrupt_pending_and_not_masked_out
    return (Config3_VEIC && IntCtl_VS && Cause_IV && !Status_BEV)
           ? Cause_RIPL > Status_IPL : Cause_IP & Status_IM;
endfunction


function interrupts_enabled
    return Status_IE && !Status_EXL && !Status_ERL && !Debug_DM;
endfunction

function implementation_dependent_wake_event
    <return true if implementation dependent waking-up event occurs>
endfunction

```

Exceptions:

Coprocessor Unusable Exception



Format: WRPGPR *rt*, *rs*

microMIPS

Purpose: Write to GPR in Previous Shadow Set

To move the contents of a current GPR to a GPR in the previous shadow set.

Description: $SGPR[SRSCtl_{PSS}, rt] \leftarrow GPR[rs]$

The contents of the current GPR *rs* is moved to the shadow GPR register specified by *SRSCtl_{PSS}* (signifying the previous shadow set number) and *rt* (specifying the register number within that set).

Restrictions:


In implementations prior to Release 2 of the Architecture, this instruction resulted in a Reserved Instruction exception.

Operation:

$SGPR[SRSCtl_{PSS}, rt] \leftarrow GPR[rs]$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: WSBH *rt*, *rs*

microMIPS

Purpose: Word Swap Bytes Within Halfwords

To swap the bytes within each halfword of GPR *rs* and store the value into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{SwapBytesWithinHalfwords}(GPR[rs])$

Within each halfword of the lower word of GPR *rs* the bytes are swapped, the result is sign-extended, and stored in GPR *rt*.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction exception.

If GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is **UNPREDICTABLE**.

Operation:

```
if NotWordValue(GPR[rt]) then
    UNPREDICTABLE
endif
GPR[rt] ← sign_extend(GPR[rs]23..16 || GPR[rs]31..24 || GPR[rs]7..0 || GPR[rs]15..8)
```

Exceptions:

Reserved Instruction

Programming Notes:


The WSBH instruction can be used to convert halfword and word data of one endianness to another endianness. The endianness of a word value can be converted using the following sequence:

```
lw    t0, 0(a1)          /* Read word value */
wsbh  t0, t0             /* Convert endiannes of the halfwords */
rotr  t0, t0, 16         /* Swap the halfwords within the words */
```

Combined with SEH and SRA, two contiguous halfwords can be loaded from memory, have their endianness converted, and be sign-extended into two word values in four instructions. For example:

```
lw    t0, 0(a1)          /* Read two contiguous halfwords */
wsbh  t0, t0             /* Convert endiannes of the halfwords */
seh   t1, t0             /* t1 = lower halfword sign-extended to word */
sra   t0, t0, 16         /* t0 = upper halfword sign-extended to word */
```

Zero-extended words can be created by changing the SEH and SRA instructions to ANDI and SRL instructions, respectively.



Format: XOR rd, rs, rt

microMIPS

Purpose: Exclusive OR

To do a bitwise logical Exclusive OR.

Description: $GPR[rd] \leftarrow GPR[rs] \text{ XOR } GPR[rt]$

Combine the contents of GPR *rs* and GPR *rt* in a bitwise logical Exclusive OR operation and place the result into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ xor } GPR[rt]$

Exceptions:

None



Format: XORI *rt*, *rs*, *immediate*

microMIPS

Purpose: Exclusive OR Immediate

To do a bitwise logical Exclusive OR with a constant.

Description: $GPR[rt] \leftarrow GPR[rs] \text{ XOR } immediate$

Combine the contents of GPR *rs* and the 16-bit zero-extended *immediate* in a bitwise logical Exclusive OR operation and place the result into GPR *rt*.

Restrictions:

None

Operation:

$GPR[rt] \leftarrow GPR[rs] \text{ xor } zero_extend(immediate)$

Exceptions:


None

Opcode Map

This chapter defines the bit-level encoding of all microMIPS64 instructions, using a series of opcode tables. The basic format of the tables is shown in [Figure 11.1](#). The topmost row contains the high-order opcode bits (in the example table shown here, bits 31..29), and the left-most column of the table lists the next most-significant bits of the opcode field (bits 28..26). Decimal and binary values are shown for both rows and columns.

An instruction's encoding is the value at the intersection of a row and column. For example, the opcode value for the instruction EX1 is 33 (decimal) or 011011 (binary). Similarly, the *opcode* value for EX2 is 64 (decimal), or 110100 (binary).

Figure 11.1 Sample Bit Encoding Table



11.1 Major Opcodes

[Table 11.2](#) defines the major opcode for each instruction. The symbols used in the table are described in [Table 11.1](#).

Every major opcode name starting with “POOL” requires a minor opcode, as defined in [Section 11.2 “Minor Opcodes”](#). All other major opcodes refer to a particular instruction.

Release 6 introduces additional nomenclature to the opcode tables for Release 6 instructions. For new instructions, bits 31:26 are generically named POP XY where X is the row number, and Y is the column number. This convention is extended to sub-opcode tables, except bits 5:0 are generically named SOP XY where X is the row number, and Y is the column number. This naming convention is applied where a specific encoded value are shared by multiple instructions.

In the opcode tables, MSB denotes either bit 15 or 31, depending on instruction size.

Table 11.1 Symbols Used in the Instruction Encoding Tables

Symbol	Meaning
*	Operation or field codes marked with this symbol are reserved for future use. Executing such an instruction must cause a Reserved Instruction Exception.
δ	(Also <i>italic</i> field name.) Operation or field codes marked with this symbol denotes a field class. The instruction word must be further decoded by examining additional tables that show values for another instruction field.
β	Operation or field codes marked with this symbol represent a valid encoding for a higher-order MIPS ISA level or a new revision of the Architecture. Executing such an instruction must cause a Reserved Instruction Exception.
\perp	Operation or field codes marked with this symbol represent instructions which are not legal if the processor is configured to be backward compatible with MIPS32 processors. If the processor is executing with 64-bit operations enabled, execution proceeds normally. In other cases, executing such an instruction must cause a Reserved Instruction Exception (non-coprocessor encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).
∇	Operation or field codes marked with this symbol represent instructions which were only legal if 64-bit operations were enabled on implementations of Release 1 of the Architecture. In Release 2 of the architecture, operation or field codes marked with this symbol represent instructions which are legal if 64-bit floating point operations are enabled. In other cases, executing such an instruction must cause a Reserved Instruction Exception (non-coprocessor encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).
Δ	Instructions formerly marked ∇ in some earlier versions of manuals, corrected and marked Δ in revision 5.03. Legal on MIPS64r1 but not MIPS32r1; in release 2 and above, legal in both MIPS64 and MIPS32, in particular even when running in “32-bit FPU Register File mode”, FR=0, as well as FR=1.
θ	Operation or field codes marked with this symbol are available to licensed MIPS partners. To avoid multiple conflicting instruction definitions, MIPS Technologies will assist the partner in selecting appropriate encodings if requested by the partner. The partner is not required to consult with MIPS Technologies when one of these encodings is used. If no instruction is encoded with this value, executing such an instruction must cause a Reserved Instruction Exception (<i>SPECIAL2</i> encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).

Table 11.1 Symbols Used in the Instruction Encoding Tables (Continued)

Symbol	Meaning
σ	Field codes marked with this symbol represent an EJTAG support instruction and implementation of this encoding is optional for each implementation. If the encoding is not implemented, executing such an instruction must cause a Reserved Instruction Exception. If the encoding is implemented, it must match the instruction encoding as shown in the table.
ε	Operation or field codes marked with this symbol are reserved for MIPS Application-Specific Extensions. If the ASE is not implemented, executing such an instruction must cause a Reserved Instruction Exception.

Table 11.2 microMIPS64 Encoding of Major Opcode Field

Major		MSB..MSB-2							
		0	1	2	3	4	5	6	7
MSB-3..MSB-5		000	001	010	011	100	101	110	111
0	000	POOL32A δ	POOL32B δ	POOL32I δ	POOL32C δ	BEQZC/JIC	BNEZC/JIALC	BLEZALC/ BGEZALC/ BGEUC	BGTZALC/ BLTZALC/ BLTUC
1	001	POOL16A δ	POOL16B δ	POOL16C δ	LWGP16	*	*	*	*
2	010	LBU16	LHU16	LWSP16	LW16	SB16	SH16	SWSP16	SW16
3	011	MOVE16	ANDI16	POOL16D δ	POOL16E δ	BEQZC16	BNEZC16	BC16	LI16
4	100	AUI32	ADDIU32	ORI32	XORI32	SLTI32	SLTIU32	ANDI32	DAUI
5	101	LBU32	LHU32	POOL32F δ	BOVC/ BEQZALC/ BEQC	BC	BALC	BGTZC/ BLTZC/ BLTC	BLEZC/ BGEZC/ BGEC
6	110	SB32	SH32	*	PCREL	SWC132	SDC132	SD32 \perp	SW32
7	111	LB32	LH32	DADDIU32 \perp	BNVC/ BNEZALC/ BNEC	LWC132	LDC132	LD32 \perp	LW32

Examples:

1. The 32-bit instruction LW32 is assigned to the major opcode LW32 with the encoding “111111”.
2. The 16-bit instruction SUBU16 is assigned to the major opcode POOL16A with the encoding “000001”.

11.2 Minor Opcodes

While major opcodes have a fixed length of 6 bits, minor opcodes are variable in length. The minor opcodes are defined by opcode tables of one, two, or three dimensions, depending on the size of the opcode. Minor opcodes less than four bits are represented in a one-dimensional table (see [Table 11.16](#)), from four to six bits in a two-dimensional table (shown in [Figure 11.1](#) and [Table 11.10](#)), and from 7 to 10 bits in a three-dimensional table ([Table 11.4](#)). In a three-dimensional table, the two-dimensional table is expanded to include a column on the right side that encodes the extra bits. In the case of minor opcodes requiring multiple table cells, the instruction name appears in all cells, but the additional entries have a black background to indicate that this opcode is blocked (see [Table 11.4](#) and the legend shown in [Table 11.3](#)).

Example:

```
SRL r1, r1, 7    binary opcode fields: 000000 00001 00001 00111 00001 000000
interpretation:  POOL32A  r1      r1      7      SRL
```

Opcode Map

hex representation: 0021 3840

All minor opcode fields are right-aligned except those in 16-bit instructions and in 32-bit instructions with a 16-bit immediate field. These left-aligned fields are defined in a bit-reverse order, which is why, in order to accommodate the variable length of the field to the right, a given row and column in POOL32I represents bit 20..22 and 23..25 instead of bit 22..20 and 25..23.

If table entries are marked grey, then not all available bits of the instruction have been used for the encoding, leaving a field of empty bits. The empty bits are shown in the instruction tables in [Chapter 5, “microMIPS Instruction Set”](#) on page 59.

Table 11.3 Legend for Minor Opcode Tables

Symbol	Meaning
OPCODE	Occupied by Opcode
OPCODE	Space Utilized by another Opcode

Table 11.4 POOL32A Encoding of Minor Opcode Field

Minor		bit 5..3									
		0	1	2	3	4	5	6	7		
bit 2..0		000	001	010	011	100	101	110	111		
										bit 9..6	
0	000	SLL32	*	SLLV	MUL	*	*	*	*	0000	0
0	000	SRL32	*	SRLV	MUH	*	*	*	*	0001	1
0	000	SRA	*	SRAV	MULU	*	*	*	*	0010	2
0	000	ROTR	*	ROTRV	MUHU	*	*	*	*	0011	3
0	000	*	*	ADD	DIV	*	*	*	*	0100	4
0	000	SELEQZ	*	ADDU32	MOD	*	*	*	*	0101	5
0	000	SELNEZ	*	SUB	DIVU	*	*	*	*	0110	6
0	000	RDHWR	*	SUBU32	MODU	*	*	*	*	0111	7
0	000	*	*	*	*	*	*	*	*	1000	8
0	000	*	*	AND	*	*	*	*	*	1001	9
0	000	*	*	OR32	*	*	*	*	*	1010	a
0	000	*	*	NOR	*	*	*	*	*	1011	b
0	000	*	*	XOR32	*	*	*	*	*	1100	c
0	000	*	*	SLT	*	*	*	*	*	1101	d
0	000	*	*	SLTU	*	*	*	*	*	1110	e
0	000	*	*	*	*	*	*	*	*	1111	f
1	001	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0	SPECIAL2 0		

Table 11.4 POOL32A Encoding of Minor Opcode Field (Continued)

2	010	COP2 0	COP2 0	COP2 0	COP2 0	COP2 0	COP2 0	COP2 0	COP2 0
3	011	UDI 0	UDI 0	UDI 0	UDI 0	UDI 0	UDI 0	UDI 0	UDI 0
4	100	*	INS	*	*	*	EXT	*	POOL32Axf 0
5	101	ε	ε	ε	ε	ε	ε	ε	ε
6	110	ε	ε		POOL32P				
7	111	BREAK32	LSA	*	ALIGN	ε	*	*	SIGRIE

Not Shown

SLL r0, r0, r0 = NOP
SLL r0, r0, 1 = SSNOP
SLL r0, r0, 3 = EHB
SLL, r0, r0, 5 = PAUSE

Table 11.5 POOL32P Encoding of Minor Extension Field

Extension		<i>bit 11..9</i>							
<i>bit 8..6</i>		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	MFC0	MTHC0		
4	100	*	*	*	*	*	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

Table 11.6 POOL32Axf Encoding of Minor Opcode Extension Field

Extension	bit 11..9										
bit 8..6	0	1	2	3	4	5	6	7			
	000	001	010	011	100	101	110	111			
0	000	TEQ	TGE	TGEU	*	TLT	TLTU	TNE	*		
1	001	ε	ε	*	ε	ε	ε	*	ε		
2	010	ε	ε	ε	ε	ε	ε	ε	ε		
3	011	MFC0	MTC0	*	*	MFC0	MTC0				
bit15..12											
4	100	ε	ε	*	*	*	BITSWAP	*	JALRC	0000	0
4	100	ε	ε	*	*	*	*	*	JALRC.HB	0001	1

Table 11.6 POOL32Axf Encoding of Minor Opcode Extension Field (Continued)

4	100	ε	*	*	*	*	SEB	*	*	0010	2
4	100	ε	*	*	*	*	SEH	*	*	0011	3
4	100	ε	*	*	*	*	CLO	MFC2	*	0100	4
4	100	ε	*	*	*	*	CLZ	MTC2	*	0101	5
4	100	ε	*	*	*	*	*	DMFC2	*	0110	6
4	100	ε	ε	*	*	*	WSBH	DMTC2	*	0111	7
4	100		*	*	*	*	*	MFHC2	*	1000	8
4	100	ε	ε	*	*	*	*	MTHC2	*	1001	9
4	100		*	*	*	*	*	*	*	1010	a
4	100	ε	ε	*	*	*	*	*	*	1011	b
4	100	*	*	*	*	*	*	CFC2	*	1100	c
4	100	ε	ε	*	*	*	*	CTC2	*	1101	d
4	100	*	*	*	*	*	*	*	*	1110	e
4	100	ε	*	*	*	*	*	*	*	1111	f

bit15..12

5	101	*	TLBP	ε	*	*	*	*	*	0000	0
5	101	*	TLBR	ε	*	*	*	*	*	0001	1
5	101	*	TLBWI	ε	*	*	*	*	*	0010	2
5	101	*	TLBWR	ε	*	*	*	*	*	0011	3
5	101	*	*	*	DI	*	*	*	*	0100	4
5	101	*	*	*	EI	*	*	*	*	0101	5
5	101	*	*	*	*	*	SYNC	*	*	0110	6
5	101	*	*	*	*	*	*	*	*	0111	7
5	101	*	*	*	*	*	SYSCALL	*	*	1000	8
5	101	*	WAIT	*	*	*	*	*	*	1001	9
5	101	*	*	*	*	*	*	*	*	1010	a
5	101	*	*	*	*	*	*	*	*	1011	b
5	101	*	*	*	*	*	*	*	*	1100	c
5	101	*	ε	*	*	*	SDBBP	*	*	1101	d
5	101	RDPGPR	DERET	*	*	*	*	*	*	1110	e
5	101	WRPGPR	ERET	*	*	*	*	*	*	1111	f

6	110	ε	ε	*	*	ε	*	*	*
---	-----	---	---	---	---	---	---	---	---

7	111	ε	ε	ε	*	*	*	*	*
---	-----	---	---	---	---	---	---	---	---

Not Shown: JR = JALR r0

Table 11.7 POOL32F Encoding of Minor Opcode Field

Minor		bit 5..3										
bit 2..0		0	1	2	3	4	5	6	7			
		000	001	010	011	100	101	110	111			
											bit 8..6	
0	000	*	*	*	ε	RINT.fmt	*	ADD.fmt	SELEQZ.fmt	000	0	
0	000	*	*	*	ε	CLASS.fmt	*	SUB.fmt	SELNEZ.fmt	001	1	
0	000	*	*	*	ε	*	*	MUL.fmt	SEL.fmt	010	2	
0	000	*	*	*	ε	*	*	DIV.fmt	*	011	3	
0	000	*	*	*		*	*	ADD.fmt	*	100	4	
0	000	*	*	*		*	*	SUB.fmt	*	101	5	
0	000	CVT.PS.S ▽	*	*	*	*	*	MUL.fmt	MADDF.fmt	110	6	
0	000	*	*	*	*	*	*	DIV.fmt	MSUBF.fmt	111	7	
1	001	*	*	*	*	*	*	*	*			
2	010	*	*	*	*	*	*	*	*			
3	011	MIN.fmt	MAX.fmt	*	*	*	MAXA.fmt	*	POOL32Fxf δ			
4	100	*	*	*	ε	*	*	*	*			
5	101	COMP.cond.S	*	CMP.cond.D	*	*	*	*	*			
6	110	*	*	*	*	*	*	*	*			
7	110	*	*	*	*	*	*	*	*			

Table 11.8 POOL32Fxf Encoding of Minor Opcode Extension Field

Extension	bit10..8										
bit 7..6	0	1	2	3	4	5	6	7			
	000	001	010	011	100	101	110	111			
bit 13..11											
0	00	*	CVT.L.fmt ▽	RSQRT.fmt Δ	FLOOR.L.fmt ▽	*	*	*	ε	000	0
0	00	*	CVT.W.fmt	SQRT.fmt	FLOOR.W.fmt	*	*	*	ε	001	1
0	00	CFC1	*	RECIP.fmt Δ	CEIL.L.fmt ▽	*	*	*	*	010	2
0	00	CTC1	*	*	CEIL.W.fmt	*	*	*	*	011	3
0	00	MFC1	*	*	TRUNC.L.fmt ▽	DMFC1	*		*	100	4
0	00	MTC1	*	*	TRUNC.W.fmt	DMTC1	*	*	*	101	5

Table 11.8 POOL32Fxf Encoding of Minor Opcode Extension Field (Continued)

0	00	MFHC1 ▽	*	*	ROUND.L.fmt ▽		*	*	*	110	6
0	00	MTHC1 ▽	*	*	ROUND.W.fmt	*	*	*	*	111	7
<i>bit 12..11</i>											
1	01	MOV.fmt	*	*	ABS.fmt	*	*	*	ε	00	0
1	01	*	*	*	NEG.fmt	*	*	*	*	01	1
1	01	*	*	*	CVT.D.fmt	*	*	*	ε	10	2
1	01	*	*	*	CVT.S.fmt	*	*	*	*	11	3
*											
2	10	*		*	*	*	*	*	*		
3	11	*	*	*	*	*	*	*	*		

Table 11.9 POOL32B Encoding of Minor Opcode Field

Minor		<i>bit 15</i>	
<i>bit 14..12</i>		0	1
		0	1
0	000	LWC2	SWC2
1	001	LWP	SWP
2	010	LDC2 ⊥	SDC2 ⊥
3	011	ε	ε
4	100	LDP ⊥	SDP ⊥
5	101	LWM32	SWM32
6	110	CACHE	*
7	111	LDM ⊥	SDM ⊥

Table 11.10 POOL32C Encoding of Minor Opcode Field

Minor		<i>bit 15</i>	
<i>bit 14..12</i>		0	1
		0	1
0	000	*	*
1	001	*	*
2	010	PREF	ST-EVA δ
3	011	LL	SC
4	100	*	*
5	101	LLDP	SCDP
6	110	LD-EVA δ	LWU
7	111	LLD	SCD

Table 11.11 LD-EVA Encoding of Minor Opcode Field

Minor		
<i>bit 11..9</i>		
0	000	LBUE
1	001	LHUE
2	010	LLWPE
3	011	*
4	100	LBE
5	101	LHE
6	110	LLE
7	111	LWE

Table 11.12 ST-EVA Encoding of Minor Opcode Field

Minor		
<i>bit 11..9</i>		
0	000	SCWPE
1	001	*
2	010	PREFE
3	011	CACHEE
4	100	SBE
5	101	SHE
6	110	SCE
7	111	SWE

Table 11.13 POOL32I Encoding of Minor Opcode Field

Minor		<i>bit 22..21</i>			
<i>bit 25..23</i>	0	0	1	2	3
	00	01	10	11	
0	000	BNZ.df	BNZ.df	BNZ.df	BNZ.df
1	001	BZ.df	BZ.df	BZ.df	BZ.df
2	010	BC1EQZ	BC1NEZC	BC2EQZC	BC2NEZC

Table 11.13 POOL32I Encoding of Minor Opcode Field

3	011	SYNCI	*	ε	ε
4	100	DATI*	DAHI*	BNZ.V	BZ.V
5	101	*	*	*	*
6	110	*	*	ε	ε
7	111	*	*	*	*
7	111	ε	ε	ε	ε

bit16

Table 11.14 POOL32S Encoding of Minor Opcode Field

Minor		bit 5..3							
bit 2..0	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	000	DSLL	DSLL32	DSLLV	DMUL *	*	*	*	000 0
0	000	DSRL	DSRL32	DSRLV	DMUH *	*	*	*	001 1
0	000	DSRA	DSRA32	DSRAV	DMULU *	*	*	*	010 2
0	000	DROTR	DROTR32	DROTRV	DMUHU *	*	*	*	011 3
0	000	DLSA *	*	DADD *	DDIV *	*	*	*	100 4
0	000	*	*	DADDU *	DMOD *	*	*	*	101 5
0	000	*	*	DSUB	DDIVU *	*	*	*	110 6
0	000	*	*	DSUBU	DMODU *	*	*	*	111 7
1	001	ε	ε	ε	ε	ε	ε	ε	
2	010	ε	ε	ε	ε	ε	ε	ε	
3	011	*	*	ε	ε	ε	ε	ε	
4	100	DINSM	DINS	DEXTU	DALIGN	DEXTM	DEXT	DINSU	POOL32Sxf δ
5	101	ε	ε	ε	ε	ε	*	ε	ε
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

bit 8..6

Table 11.15 POOL32Sxf Encoding of Minor Opcode Extension Field

Extension		bit 11..9									
bit 8..6	0	0	1	2	3	4	5	6	7		
	000	001	010	011	100	101	110	111			
0	000	*	*	*	*	ε	ε	ε	ε		
1	001	*	ε	ε	ε	ε	ε	ε	ε		
2	010	ε	ε	ε	ε	ε	ε	ε	ε		
3	011	DMFC0	DMTC0	*	*	DMFC0	DMTC0	*	*		
bit15..12											
4	100	ε	ε	*	*	*	DBITSWAP *	*	*	0000	0
4	100	ε	ε	*	*	*	*	*	*	0001	1
4	100	ε	ε	*	*	*	*	*	*	0010	2
4	100	*	*	*	*	*	*	*	*	0011	3
4	100	ε	*	*	*	*	DCLO	*	*	0100	4
4	100	ε	ε	ε	*	*	DCLZ	*	*	0101	5
4	100	ε	ε	ε	*	*	*		*	0110	6
4	100	ε	ε	*	*	*	DSBH		*	0111	7
4	100	*	*	*	*	*	*	*	*	1000	8
4	100	ε	ε	*	*	*	*	*	*	1001	9
4	100	*	*	*	*	*	*	*	*	1010	a
4	100	ε	ε	*	*	*	*	*	*	1011	b
4	100	*	*	*	*	*	*	*	*	1100	c
4	100	ε	ε	*	*	*	*	*	*	1101	d
4	100	*	*	*	*	*	*	*	*	1110	e
4	100	ε	*	*	*	*	DSHD	*	*	1111	f
5	101	*	*	*	*	*	*	*	*		
6	110	ε	ε	*	*	*	*	*	*		
7	111	ε	ε	ε	ε	*	*	*	*		

Table 11.16 POOL16A Encoding of Minor Opcode Field

Minor	
bit 0	
0	ADDU16
1	SUBU16

Table 11.17 POOL16B Encoding of Minor Opcode Field

Minor	
<i>bit 0</i>	
0	SLL16
1	SRL16

Table 11.18 POOL16C Encoding of Minor Opcode Field

Minor		<i>bit 6..4</i>							
		0	1	2	3	4	5	6	7
<i>bit 9..7</i>		000	001	010	011	100	101	110	111
0	000	NOT16	XOR16	NOT16	XOR16	NOT16	XOR16	NOT16	XOR16
1	001	AND16	OR16	AND16	OR16	AND16	OR15	AND16	OR16
2	010	LWM16	SWM16	LWM16	SWM16	LWM16	SWM16	LWM16	SWM16
3	011	JRC16	JALRC16	JRCADDIUSP	BREAK16	JRC16	JALRC16	JRCADDIUSP	SDBBP16
4	100	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP
5	101	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP
6	110	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP
7	111	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP	MOVEP

For Release 6:

- NOT16, AND16, XOR16, OR16, LWM16, SWM16, and BREAK16, and SDBBP16 instructions have been repositioned in POOL16C and are not shown in the above table.
- JRADDIUSP has been converted to JRCADDIUSP and repositioned in POOL16C.
- MOVEP has moved from POOL16F to POOL16C.

Table 11.19 POOL16D Encoding of Minor Opcode Field

Minor	
<i>bit 0</i>	
0	ADDIUS5
1	ADDIUSP

Table 11.20 POOL16E Encoding of Minor Opcode Field

Minor	
<i>bit 0</i>	
0	ADDIUR2
1	ADDIUR1SP

Table 11.21 PCREL Encoding of Minor Opcode Field

Extension		<i>bit 20..18</i>							
<i>bit 17..16</i>		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00	ADDIUPC	ADDIUPC	LWPC	LWPC	LWUPC	LWUPC	LDPC	*
1	01	ADDIUPC	ADDIUPC	LWPC	LWPC	LWUPC	LWUPC	LDPC	*
2	10	ADDIUPC	ADDIUPC	LWPC	LWPC	LWUPC	LWUPC	LDPC	AUIPC
3	11	ADDIUPC	ADDIUPC	LWPC	LWPC	LWUPC	LWUPC	LDPC	ALUIPC

11.3 Floating Point Unit Instruction Format Encodings

Instruction format encodings for the floating point unit are presented in this section.

If the instruction allows Single, Double and Pair-Single formats, the following encoding is used:

Table 11.22 Floating Point Unit Format Encodings - S, D, PS

<i>fmt field</i>		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
0	0	S	Single	32	Floating Point
1	1	D	Double	64	Floating Point
2	2	PS	Paired Single	2 × 32	Floating Point
3	3	Reserved for future use by the architecture.			

If the instruction only allows Single and Double formats, the following encoding is used:

Table 11.23 Floating Point Unit Format Encodings - S, D 1-bit

<i>fmt field</i>		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
0	0	S	Single	32	Floating Point

Table 11.23 Floating Point Unit Format Encodings - S, D 1-bit

<i>fmt</i> field		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
1	1	D	Double	64	Floating Point

Table 11.24 Floating Point Unit Instruction Format Encodings - S, D 2-bits

<i>fmt</i> field		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
0	0	S	Single	32	Floating Point
1	1	D	Double	64	Floating Point
2, 3	2, 3	Reserved for future use by the architecture.			

If the instruction allows Single, Word and Long formats, the following encoding is used:

Table 11.25 Floating Point Unit Format Encodings - S, W, L

<i>fmt</i> field		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
0	0	S	Single	32	Floating Point
1	1	W	Word	32	Integer
2	2	L	Long	64	Integer
3	3	Reserved for future use by the architecture.			

If the instruction allows Double, Word and Long formats, the following encoding is used:

Table 11.26 Floating Point Unit Format Encodings - D, W, L

<i>fmt</i> field		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
0	0	D	Double	64	Floating Point
1	1	W	Word	32	Integer
2	2	L	Long	64	Integer

Table 11.26 Floating Point Unit Format Encodings - D, W, L

<i>fmt</i> field		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex				
3	3	Reserved for future use by the architecture.			

Compatibility

This chapter covers various aspects of compatibility. microMIPS64 is the preferred replacement for the existing MIPS16e ASE and uses the same mode-switch mechanism. Although microMIPS includes almost all MIPS64 instructions and therefore does not require the original MIPS64 encodings, initially it will be implemented together with MIPS64-encoded instruction execution.

12.1 Assembly-Level Compatibility

microMIPS64 includes a re-encoding of the MIPS64 instructions, including all ASEs and UDI space. Therefore, microMIPS provides assembly-level compatibility. Only the following cases cause some side effects:

- **Re-encoded MIPS64 instructions with reduced operand fields**

There are 3 classes of reduced fields:

1. *Reserved or unsupported bits and encodings.* This category is not a problem because utilizing a reserved or unsupported field causes an exception, no operation, or undefined behavior, and often these cannot be accessed by the compiler anyway. An example of this category is the ‘fmt’ field.
2. *Bit fields and ranges which are defined but typically never used.* This category is usually not a problem. The assembler generates an error message if a constant is outside of the re-defined range.
3. *Bit fields which are used but were reduced in order to utilize the new opcode map most efficiently.* The handling of these cases is similar to category 2 above—compilers do not generate such scenarios, and assemblers generate error messages. In the latter case, the programmer has to either fix the code or switch to the MIPS64 encoding.

- **Re-encoded Branch and Jump instructions**

Branch instructions support 16-bit aligned branch target addresses, providing full flexibility for microMIPS. Because the offset field size of the 32-bit encoded branch instructions is the same as the MIPS64-encoded instructions, and because all branch target addresses of the MIPS64 encoding are 32-bit aligned, the branch range in microMIPS is smaller. This is partially compensated by the smaller code size of microMIPS.

Jump instructions also support 16-bit aligned target addresses. This reduces the addressable target region for J, JAL to 128 MB instead of 256 MB. For these instructions, the effective target address is in the ‘current’ 128 MB-aligned region. For larger ranges, the jump register instructions (JR, JRC, and JRADDIUSP) can be used.

- **MIPS64 assembly instructions manually encoded using the .WORD directive**

Manual encoding of MIPS64 assembly instructions can be used in assembly code as well as assembly macros in C functions. To differentiate between microMIPS-encoded instructions and other encoded instructions or data, the following compiler directives have been introduced:

```
.set micromips    ; instruction stream is microMIPS

.set nomicromips ; instruction stream is MIPS64

.insn            ; If in microMIPS instruction stream mode, the location associated
                ; with the previous label is aligned to 16-bit bits instead of
                ; 32-bits
                ; If in microMIPS instruction stream mode and if the previous
                ; label is loaded to a register as the target of a jump or branch,
                ; the ISAMode bit is set within the branch/jump register value.
```

The programmer must use these directives to encode instructions in microMIPS.

For example, to manually encode a microMIPS NOP:

```
.set micromips

label1:
    .insn
    .word 0          ; label1 location - represents microMIPS NOP32 instruction
label2:
    .insn
    .half 0x0c00 ; label2 location - represents microMIPS NOP16 instruction
label3:
    .half 0x0c00 ; label3 location - represents data value of 3072 (decimal)
```

To manually encode a MIPS64 NOP:

```
.set nomicromips
.word 0          ; represents MIPS64 NOP instruction
```

For MIPS64 instruction stream mode, the “.insn” directive has no effect.

- **Branch likely instructions**

microMIPS does not support branch likely instructions in hardware. Assembly-level compatibility is maintained because assemblers replace branch likely instructions either by an instruction sequence or by a regular branch instruction, and they perform some instruction reordering if reordering is possible.

12.2 ABI Compatibility

microMIPS is compatible with the existing ABIs o32, n32, and n64 calling conventions. However, a few new relocation types need to be added to these ABIs for microMIPS support, as some of the additional offset field sizes required for microMIPS become visible to the linker. For example, the offset fields of J and SW using GP are visible to the linker, while B and SWSP are hidden within the object files.

Functions remain 32-bit aligned as in the MIPS64 encoding as well as MIPS16e. This guarantees that static and dynamic linking processes can link microMIPS object files with MIPS64 object files.

Programs can be composed of both microMIPS and MIPS64 modules, using either the JALX instructions (and/or JR instructions with setting the ISAMode bit appropriately) to switch instruction set modes when calling routines compiled in an ISA different from that of the caller routine.

microMIPS provides flexibility for potential future ABIs.

12.3 Branch and Jump Offsets

microMIPS branch targets are half-word (16-bit) aligned to match half-word sized instructions. Please refer to [Section 2.5, "Branch and Jump Offsets."](#)

12.4 Relocation Types

Compiler and linker toolchains need to be modified with new relocation types to support microMIPS. Reasons for these new relocation types include:

1. The placement of instruction halfwords is determined by memory endianness. MIPS64 instructions are always of word size, so there were no halfword placement issues.
2. microMIPS has 7-bit, 10-bit and 16-bit PC-relative offsets.
3. Branch and Jump offset fields are left-shifted by 1 bit (instead of 2 bits in MIPS64) to create effective target addresses.
4. Some code-size optimizations can only be done at link time instead of compile time. Some new relocation types are used solely within the linker to keep track of address and data information.

12.5 Boot-up Code shared between microMIPS64 and MIPS64

In some systems, it would be advantageous to place both microMIPS64 and MIPS64 executables in the same boot memory. In that way, a single system could be used for either instruction set.

To enable this, a binary code sequence is required that can be run in either instruction set and change code paths depending on the instruction set that is being used.

The following binary sequence achieves this goal:

```
0x1000wxyz // where w,x,y,z represent hexadecimal digits
0x00000000
```

For the MIPS64 instruction set, this binary sequence is interpreted as:

```
BEQ $0, $0, wxyz // branch to location of more MIPS64 instructions
NOP
```

For the microMIPS instruction set, this binary sequence is interpreted as:

```
ADDI32 $0, $0, wxyz // do nothing
NOP                // fall through to more microMIPS instructions
```

12.6 Coprocessor Unusable Behavior

When a coprocessor instruction is executed when the associated coprocessor has not been implemented, it is allowed for the RI exception to be signalled instead of the Coprocessor Unusable exception. Please refer to [Section 2.6](#), "Coprocessor Unusable Behavior."

12.7 Other Issues Affecting Software and Compatibility

microMIPS instructions can cross cache lines and page boundaries. Hardware must handle these cases so that software need not avoid them. Since MIPS64 requires instructions to be 32-bit aligned, there is no forward compatibility issue when transitioning to microMIPS.

References

This appendix lists other publications available from Imagination Technologies, some of which are referenced elsewhere in this document. They may be included in the `$MIPS_HOME/$MIPS_CORE/doc` area of a typical soft or hard core release, or in some cases may be available on the MIPS web site, <http://www.imgtec.com>.

- MIPS® Architecture For Programmers, Volume I: Introduction to the MIPS64® Architecture
- MIPS® Architecture For Programmers, Volume II: The MIPS64® Instruction Set
- MIPS® Architecture For Programmers, Volume III: The MIPS64® and microMIPS64™ Privileged Resource Architecture

References

Revision History

Revision	Date	Description
1.08	November 25, 2009	<ul style="list-style-type: none"> Clean-up for external release.
1.09	January 7, 2010	<ul style="list-style-type: none"> Added shared boot-up code sequence in Compatibility Chapter.
3.00	March 25, 2010	<ul style="list-style-type: none"> Changed document revision numbering to match other Release 3 documents. Hopefully this will be less confusing. Moved MIPS32/64 version of JALX to Volume II-A.
3.01	October 30, 2010	<ul style="list-style-type: none"> User mode instructions not allowed to product UNDEFINED results. Updated copyright page. Removed Margin Note - “Preliminary - Subject to Change” in some chapters.
3.02	December 6, 2010	<ul style="list-style-type: none"> POOL32Sxf binary encoding was incorrect for individual instruction description pages.
3.03	December 10, 2010	<ul style="list-style-type: none"> microMIPS AFP versions security reclassification.
3.04	March 21, 2011	<ul style="list-style-type: none"> RSQRT/RECIP does not need 64-bit FPU. MADD.fmt/NMADD.fmt/MSUB.fmt/NMSUB.fmt psuedo-code was incorrect for PS format check.
3.05	April 4, 2011	<ul style="list-style-type: none"> The text description was incorrect for the offset sizes for these instructions - CACHE, LDC2, LL, LWC2, LWL, LWR, PREF, SDC2, SWL, SWR. CACHE & WAIT instruction descriptions were using the wrong instruction bit numbers. LWU was incorrectly included int the microMIPS32 version.
3.06	October 17, 2012	<ul style="list-style-type: none"> CVT.D.fmt and CVT.S.fmt were in wrong positions within Table POOL32Fxf.
3.07	October 26, 2012	<ul style="list-style-type: none"> Fix Figure 6.1 - columns & rows were transposed from the real tables.
5.00	December 14, 2012	<ul style="list-style-type: none"> Some of the microMIPS instructions were not listed in alphabetical order. Fixed. No content change. R5 changes: DSP and MT ASEs -> Modules NMADD.fmt, NMSUB.fmt - for IEEE2008 negate portion is arithmetic.
5.01	December 16, 2012	<ul style="list-style-type: none"> No technical context change: Update cover with microMIPS logo Update copyright text. Update pdf filename.
5.03	August 21, 2012	<ul style="list-style-type: none"> Resolved inconsistencies with regards to the availability of instructions in MIPS32r2: MADD.fmt family (MADD.S, MADD.D, NMADD.S, NMADD.D, MSUB.S, MSUB.D, NMSUB.S, NMSUB.D), RECIP.fmt family (RECIP.S, RECIP.D, RSQRT.S, RSQRT.D), and indexed FP loads and stores (LWXC1, LDXC1, SWXC1, SDXC1). These instructions are required to be available in all FPUs. .

Revision History

Revision	Date	Description
5.04	January 15, 2014	<p>LLSC Related Changes</p> <ul style="list-style-type: none"> Added ERETNC. New. Modified SC handling: refined, added, and elaborated cases where SC can fail or was UNPREDICTABLE. <p>XPA Related Changes</p> <ul style="list-style-type: none"> Added MTHC0, MFHC0 to access extensions. All new. Modified MTC0 for MIPS32 to zero out the extended bits which are writeable. This is to support compatibility of XPA hardware with non XPA software. In pseudo-code, added registers that are impacted. MTHC0 and MFHC0 - Added RI conditions.
6.0	February 27, 2015	<ul style="list-style-type: none"> Release 6 compatible microMIPS. See Section 2.7, "Release 6 of the MIPS Architecture," for instructions that have been added, removed, and recoded.
6.01	June 9, 2015	<ul style="list-style-type: none"> Removed the Release 6 NAL instruction; it is not required in microMIPS. Removed the “Jump and Link Restartability” paragraph from JAL-type instructions; it is not applicable for compact jumps. Fixed text in jump instructions related to the behavior of ISAMode switching, or lack thereof, in microMIPS. Removed delay-slot references; all branches/jumps are compact. Removed references to JALX. Removed LWXS (bug). MOVEP: in encoding, changed bit 2 to 1 (bug). All PC-related instructions: qualify PC with 0x3. Always word aligned. microMIPS only (ADDIUPC, LWPC, AUIPC, ALUIPC). Release 6 BC: shift-corrected to 1 bit. microMIPS only. JALRC, JALRC.HB: replace Config1.CA with Config3.ISA (bug). Added Release 5 TLBINV/TLBINVF (incorrectly excluded from book). Added Release 6 DVP/EVP instructions. Added new Release 6 LLX/SCX family instructions. General opcode map cleanup for consistency with Release 6. CACHE, PREF, LL, SC, LLD, SCD, LLX, SCX, LLDX, AND SCDX offsets changed to 9 bits. LWC2, SWC2, LDC2, AND SDC2 offsets changed to 11-bits for consistency with MIPS Release 6. <p>Specific opcode map changes:</p> <ul style="list-style-type: none"> Moved BGTZC/BLTZC/BLTC to (5,6) location Moved BLEZC/BGEZC/BGEC to (5,7) location to free up 16-bit instruction rows for 16-bit instructions (only) Moved BEQZC/JIC to (0,4) location. Moved BNEZC/JIALC to (0,5) location bit for differentiating EQ vs NE type made consistent with other branches of this type.
6.02	July 13, 2015	<ul style="list-style-type: none"> Added SIGRIE instruction. Added Config5.SBRI dependence to SDBBP16 as in SDBBP(32). Corrected mistake in 6.01. AUIPC PC should not be qualified with 0x3. Added misaligned ld/st support. ld/st that must be aligned are indicated as such. Added RDHWR with sel field - read all.

Revision	Date	Description
6.03	November 13, 2015	<p>microMIPS32 and microMIPS64:</p> <ul style="list-style-type: none"> • J/JAL now indicated as deprecated (but not removed). • DVP: added text indicating that a disabled VP will not be re-enabled for execution on deferred exception. • CACHE/CACHEE: Undefined operations are really NOP. • CMP.condn.fmt: removed fmt related text in description section. .S/.D explicitly encoded. • Fixed minor textual typos in MAXA/MINA.fmt functions. • DERET: restriction – if executed out of debug mode, then RI, not UNDEFINED. • TLBWR: Updated reference to Random. No longer supported in Release 6. • PCREL instructions: added PCREL minor opcode table, fixed conditional text bugs in register reference. • BGTZ/BLTZ/BLTC major opcodes fixed. • JALRC16: in operation section, PC+2 assigned to GPR[31], not PC+4. • LW16: offset is an unsigned value. Legacy bug. • BC16: In operation section change PC<-PC+target_offset to PC+2+target_offset. • Fixed minor encoding of MUH. • Fixed typo ROUND/TRUNC/FLOOR/CEIL.W.fmt. Range value should be $2^{31}-1$ not $2^{63}-1$. <p>microMIPS64 only:</p> <ul style="list-style-type: none"> • DMFC0/DMTC0: now indicates what happens with 32-bit COP0 registers.
6.04	June 6, 2016	<p>microMIPS32 and microMIPS64:</p> <ul style="list-style-type: none"> • RDHWR: Changed Double-Width LLX/SCX to Paired LL/SC. • DMTC2: Changed CPR[2, rd, sel] to CP2CPR[Impl]. • WAIT: Fixed a bit range typo. • LSA: Removed the word optional; the scaling shift on rs is not optional. • SYSCALL, TEQ, TGE, and TGEU: If COP0 <i>BadInstr</i> is implemented, the <i>code</i> field may be obtained from <i>BadInstr</i>. • JALRC, JALRC.HB, JIALC, and JIC: Added parentheses to condition for $PC \leftarrow temp$ in the Operation pseudocode. <p>microMIPS32:</p> <ul style="list-style-type: none"> • Removed the LLX, LLXE, SCX, and SCXE instructions. • Added the LLWP, LLWPE, SCWP, and SCWP instructions. <p>microMIPS64:</p> <ul style="list-style-type: none"> • Removed the LLDX and SCDX instructions. • Added the LLDP, LLWP, LLWPE, SCDP, SCWP, and SCWPE instructions.

Revision History

Revision	Date	Description
6.05	December 15, 2016	<p>microMIPS32 and microMIPS64:</p> <ul style="list-style-type: none">• Added CRC32B, CRC32H, CRC32W, CRC32CB, CRC32CH, CRC32CW.• DVP/EVP instructions incorrectly used ‘rs’. Changed to use ‘rt.’• Added GINVI, and GINVT instructions.• SC, SCE, SCWP, SCWPE: Updated description for uncached handling.• MTHC0: updated description, fixed typo. ‘COP2’ changed to ‘COP0’.• MTC0: changed <i>Config5_{MVH}</i> to new <i>Config5_{XPA}</i>.• DERET: updated pseudocode to describe what happens if <i>Debug_{DM}</i>=0.• TLT, TLTU, TNE: Mention that contents of the <i>code</i> field can be retrieved from COP0 BadInstr if present.• Added ArchitectureRevision(), IsCoproprocessorRegisterImplemented(), and IsCoproprocessorRegisterExtended() pseudocode descriptions (ARM only, not AFP) <p>microMIPS64:</p> <ul style="list-style-type: none">• Added CRC32D, CRC32CD.• SCD, SCDP: Updated description for uncached handling.• LLDP: Fixed typos. Swapped rd and rt in the GPR references and description; in the pseudocode, doubled the bit range in both cases where GPR is loaded. (typos).• SCWP, SCWPE: fixed typo due to conditional text.